Silverback Studios

Built to Innovate, Managed to Comply



Background

Case

Silverback Studios was founded seven years ago in Amsterdam by three Stanford alumni - Joel Schumer and Bart Dijkstra, both Dutch, and Ella Jones, American - who had spent their university years building small experimental games. All three were lifelong gamers, and what started as a shared hobby became, almost by accident, a company.

Their first real project was a post-apocalyptic survival game, *Meridian Eclipse*, built largely in evenings and weekends with a handful of friends. It had no marketing budget and no formal launch, it simply found its audience through online gaming forums. Within two years, that grassroots following had grown large enough to catch the attention of Sony Interactive Entertainment, who acquired the studio and funded a full cross-platform version of the game. The new version of Meridian Eclipse was developed over the following two years and successfully released on PlayStation, Xbox, and PC. The release became a genuine breakout hit - critically acclaimed, commercially massive, and enough to put Silverback on the map as one of the most exciting studios in Europe. A three-year development cycle for the sequel followed, and by the time it neared release, Silverback was considered one of the most closely watched studios in the industry.

Silverback's growth mirrored its rising profile. The company started with 20 people, mostly friends of the three founders, and a few early industry contacts. By the time Sony acquired the company, that number had already grown to around 650. The investment and development ramp-up that followed pushed headcount to roughly 3,000 within a year, and by the time the breakout release shipped, Silverback employed close to 7,000 people. At its peak - in the run-up to the sequel, with the studio's reputation at its highest point - headcount had reached roughly 9,500, a significant share of them contractors and freelancers brought on in the final year to help meet the release deadline.

Silverback is headquartered in Amsterdam, with a large co-development studio in Kraków supporting production at scale, and a smaller motion-capture and audio studio in Los Angeles. The studio is led jointly by its three founders, with Joel being its public face.

Three years into development on the sequel, *Meridian Eclipse II*, Bart, Joel and Ella feel that the mood inside Silverback has shifted: something about the way the studio used to work has gone missing. Silverback built its reputation in large part on its game visuals: its real-time rendering of the original game's vast apocalyptic skylines and exotic landscapes, rendered in full detail as players moved through them, with weather, erosion, and collapsing structures. It was the kind of technical innovation that got the studio noticed in the first place. For the sequel, the plan was to push that further: richer environments, more dynamic systems. Instead, progress has been grinding. The rendering team has hit the same wall for months: every attempt to extend the technique to the sequel's larger, more detailed world comes at a cost the engine can't absorb without new architecture. So far nobody has produced the kind of elegant workaround that used to appear, seemingly effortlessly, in the studio's early years. It's as if the innovative ideas that once came easily simply aren't showing up the way they used to. In addition, several leading game designers, technical artists, and art directors have left the company.

Faced with a problem they can't quite name, let alone fix, Bart, Joel, and Ella decide it's time to bring in outside help. The obvious call is Jessica "Jess" Saltz, a Stanford classmate from the same cohort as the three of them, German, who went on to found her own organizational development consultancy after several years in industry. Jess's firm has built a reputation for a distinctly evidence-based approach.

Step 1. Defining the problem

Question

Before deciding on any solution, the problem itself needs to be clearly defined. Assess whether that is currently the case. In particular, consider:

a) *What is the problem to be solved?*

b) *How clearly defined is the problem (what, who, when, where, why)?*

c) *What assumptions underlie Bart, Joel, and Ella framing of the problem?*

d) *Is it clear what the organizational consequences of the problem are?*

e) *Is it clear how serious and urgent the problem is?*

f) *Is it clear what the logic model (causal mechanism) is: what are the underlying cause(s), how do they lead to the problem, and how does the problem consequently lead to the organizational consequences?*

Answer/feedback

A. Unclear: the founders sense that "innovation/creativity has declined," evidenced by slower technical progress, but this is an impression, not a defined problem.

B. Poorly defined on nearly every dimension. *What*: vague ("things take longer," "ideas aren't coming"). *Who*: unclear — the whole studio? specific teams (rendering)? *When*: only loosely dated to the sequel's development, no clear onset. *Where*: unclear whether this is studio-wide or localized to certain functions/locations (Amsterdam vs. Kraków vs. LA). *Why*: entirely unaddressed so far.

C. That the core problem is "innovation" or "creativity", rather than, for example, work load, team fatigue, loss of key talent, or scale itself.

D. No: no data yet on cost, delay, or financial impact of the slowdown.

E. No: "taking longer than expected" is subjective; no baseline or threshold has been established for what "too long" or "too slow" actually means.

F. No: there's no articulated chain from cause to problem to consequence yet.

Step 2. Identifying evidence gaps

Case

Jess's first meeting is with Joel and Ella. After they've walked her through the backstory, she asks them several questions. What, exactly, do they think the problem is? And what evidence do they actually have that it's a real problem, rather than a feeling or a hunch?

Joel explains that Silverback runs a sophisticated and detailed data warehouse that pulls together information from its financial, operational, and HR systems into a single integrated environment. Most of this data is captured automatically, straight from the source systems. Among the core metrics is the number of technical or structural changes made to the game's architecture and design — every change is logged in the system and tied to the personnel number of whoever was responsible for it. Changes considered ‘substantial’ innovations are also filed as patents, through Sony's corporate patent process. Exactly this data, Joel explains, is what triggered their concern: over the past two years, the number of changes has been declining, and the patent numbers confirm it isn't just an artifact. The reason this matters, he explains, is because the new sequel doesn't just need to be better, faster, and more visually striking than the previous release, it also needs to keep pace with whatever the console makers (PlayStation, Xbox) are pushing forward on their end. Fall behind on either front and *Meridian Eclipse* will lose ground to competitors. And lost ground means lost market share, and lost market share means lost revenue, etc.

Next, Jess focuses on the 'who' and 'when': who exactly is developing/designing these changes? Did they find any differences between teams and departments? Did they find any correlations with other KPIs/variables? Joel explains that the people responsible for these changes are the lead programmers, the systems designers, the technical artists, and the principal rendering and graphics engineers - they are considered the technical and creative ‘core’ of the company. Ella adds they indeed found differences. Teams led by senior managers who've been at Silverback since the early days log noticeably more changes and innovations than teams led by those who joined within the last two years. In addition, when they cross-referenced the logged changes against HR records, they found a strong association with lower employee satisfaction, higher absenteeism, and higher turnover.

“That’s an interesting finding,” Jess says. “Did you also check associations with perceived workload, fatigue, staffing, training, education, experience, tenure, age, employment status, or any other HR metric?” “Yes,” says Ella, “but all these correlations were not significant and practically irrelevant.” Jess then moves to the other sources of evidence: have they consulted experienced managers, or the stakeholders themselves - the people actually responsible for the changes and innovations? Joel and Ella explain they did not - nor did they consult the research literature on innovative performance/behavior.

Question

It looks like the only evidence Joel and Ella have consulted concerns organizational data. What other evidence, from which sources, is needed to establish whether (and if so, how,

for whom, and when) Silverback has a problem with innovation? For each of the evidence sources, describe what you would like to know.

Answer/feedback

Practitioners > What's missing is evidence from experienced managers with broad exposure to innovation and creative output across teams/departments. Rather than a generic sample, the panel should be deliberately structured to compare two contrasts already suggested by the organizational data:

- **High-output vs. low-output teams**: managers leading teams with consistently high numbers of logged changes, compared with managers leading teams with consistently low numbers
- **Long tenured vs. recently joined managers**: managers who've been with Silverback since the early days, compared with managers who joined within the last two years.

Relevant questions to ask:

1. Based on your experience, what do you think could explain a decline in the number of significant technical/structural changes being logged (for example workload, team size, management style, tooling, process/approval overhead, or something else)?
2. Have you noticed differences between teams in how much innovative work gets produced, and if so, what do you think explains that difference?
3. In your experience, does a lower number of *logged* changes necessarily mean less innovation is happening, or could some innovative work be happening but not getting captured in the system?

Scientific literature > The main question worth pursuing is:

- What have empirical studies found to be the most important predictors of innovative work behavior?

A narrower, follow-up search could be conducted to see what the scientific literature says about the impact of contextual variables such as rapid organizational growth, span of control, organization size, and increased bureaucracy.

Stakeholders > The relevant stakeholders here are the lead programmers, systems designers, technical artists, and principal rendering and graphics engineers - the people directly credited with the logged changes, and the ones whose actual day-to-day experience the organizational data can't capture.

Relevant questions to ask:

1. It has been found that, over the past two years, the number of significant technical/structural changes being logged into the system has declined. Based on your experience, what do you think has contributed to this decline (for example: workload, work environment, training, rapid organizational growth, increased bureaucracy, lack of autonomy)?
2. Do you feel supported by your manager and the wider organization when proposing something new or untested?
3. Is there anything about how Silverback is currently organized (for example: team structure, approval processes) that makes it harder to innovate than it used to be?

Note: The factors listed in Q1 should be updated once the scientific literature review is complete, so the question reflects the actual predictors the research identifies — rather than a list based on assumptions.

Step 3. Acquiring the evidence

Question

In the previous step, you identified what evidence is needed from each of the three missing sources: practitioners, the scientific literature, and stakeholders. For each source, describe how you would acquire the evidence (i.e. which method you would use).

Answer/feedback

Practitioners: Given the comparative design (high-output vs. low-output teams, long-tenured vs. recently joined managers) semi-structured interviews are a better fit than a survey. The core questions ("… what do you think could explain a decline in the number of changes being logged ..." and “have you noticed differences between teams …”) probably need probing to get a real answer; a survey would likely produce succinct, generic answers.

Given the four comparison groups, the sample should include a reasonable number from each cell (enough to identify patterns, rather than individual anecdotes) preferably 3 to 4 managers per group.

Scientific evidence: We would conduct a systematic search in academic databases such as ABI/Inform, Business Source Elite/Premier, and PsycINFO. Relevant search terms are ‘innovation’, ‘innovative performance’ and ‘innovative behavior’. When time is a constraint,

a search with AI search tools such as Consensus and Elicit can be a good alternative, using the questions as formulated under Step 2.

Stakeholders: Given the scale - a large number of lead programmers, systems designers, technical artists, and rendering engineers - a survey with both open and close-ended (check all that apply) type of questions is the more realistic choice than interviews and focus groups. In this case anonymity matters more than usual: staff may worry that a candid answer ("I haven't had time to innovate because of X") could reflect badly on them (or their manager). As with the managers, the same sampling logic applies here: stakeholders from both high-output and low-output teams, and ideally from teams under both long-tenured and recently joined managers.

Step 4. Appraising the trustworthiness of the evidence

Case

Jess and her team decide that semi-structured interviews (covering the same three core questions with room for follow-up) would be the best approach. The answers will be recorded, transcribed, and then summarized with the use of an AI tool like ChatGPT. They also decide that the interviews should be conducted by someone independent, since managers are effectively being compared against each other. Confidentiality should be made explicit up front. Jess approaches Silverback's HR department to select a purposive sample of 3 to 4 managers per comparison group, so a total of 12 to 16.

Question

Do you feel that this procedure, and the resulting 12 to 16 responses, will yield trustworthy evidence? Explain why or why not, and identify any limitations.

Answer/feedback

This is a moderately trustworthy design, but with some limitations worth considering.

The strengths: independent interviewers reduce bias; the comparative sampling (high/low output × tenure) directly targets the pattern already found in the data; the semi-structured format allows probing, likely producing richer answers than a survey; confidentiality upfront should improve candor.

Limitations: 12–16 managers is a small sample - enough to spot patterns, not enough for statistical generalization; using an AI tool to summarize transcripts introduces a real risk of losing nuance or subtly reshaping what people actually said, especially for a comparative analysis where precise wording differences between groups matter; and managers, even with confidentiality promised, may still shade their answers knowing they're effectively being evaluated against peers. Thus, some social desirability bias is likely unavoidable.

Question

How would you judge the trustworthiness of the scientific evidence if it was obtained as described in Step 3?

Answer/feedback

Trustworthy, provided the process is followed as described. It's important, however, that the review specifically includes meta-analyses, controlled studies, and longitudinal studies in which the relationship between innovative behavior and predicting factors was quantitatively measured, since these designs provide stronger evidence than cross-sectional studies. In addition, the included studies should be critically appraised to determine their methodological quality, and effect sizes should be reported or calculated, preferably with confidence intervals.

Case

Jess and her team decide to run the stakeholder survey only after the scientific literature review is complete. That way, the strongest predictors of innovative behavior identified by the review can be built directly into the close-ended ("check all that apply") question asking stakeholders about possible contributing factors, rather than guessing at categories in advance. This keeps the survey targeted, instead of turning it into a fishing expedition. Jess approaches Silverback's HR department to obtain the email addresses of the target population: lead programmers, systems designers, technical artists, and rendering engineers, drawn from both high-output and low-output teams, and from teams under both long-tenured and recently joined managers. The survey is set up through SurveyMonkey, and an invitation is sent to 915 employees, through a procedure that guarantees anonymity. Of these, 693 respond, which is a response rate of about 75%.

Question

How would you judge the trustworthiness of the findings drawn from these 693 responses?

Answer/feedback

Good, but not excellent. A 75% response rate is strong and reduces non-response bias, and anonymity should improve candor on sensitive questions. However, there is one thing worth checking before treating the findings as fully trustworthy: balance across the four comparison cells. A high overall response rate doesn't guarantee even response *within* each group. If, say, staff in low-output teams responded at a different rate than those in high-output teams, the comparison itself could be skewed even with 693 responses in total.

Beyond that, trustworthiness increases if there's strong agreement or homogeneity among the answers, consistent responses across staff pointing to the same factors are harder to dismiss as noise. Finally, the real test of trustworthiness is convergence across sources: do these findings align with what the manager interviews, the scientific evidence and the organizational data show?

Step 5. Aggregating the evidence

Case

Over the following week, Jess and a senior member of her team conduct semi-structured interviews with 16 managers: 9 long tenured and 7 who joined within the past year. The pattern already suggested by the organizational data mostly holds: 7 of the 9 long-tenured managers lead teams with a high number of logged changes, and 6 of the 7 recently joined managers lead teams with a low number. But not entirely. Two of the long-tenured managers lead teams with low output, and one of the recently joined managers leads a team with high output.

Question

Download the report with the managers' individual answers <<here>> and summarize the main theme(s). What does this suggest about a possible explanation for the decline in logged changes and patents. You can use an AI tool such as Claude or ChatGPT to support you.

Answer/feedback

Two main themes emerge.

- First, managers of high-output teams consistently describe an environment where any idea can be raised, no matter how unconventional, without fear of being shot down or judged for it. They frame this as the studio's founding spirit. Managers of low-output teams describe a more transactional, efficiency-focused style: staying in scope, skepticism toward unsolicited ideas ("innovation for its own sake"), a high bar before a suggestion is considered.
- This pattern holds even where tenure doesn't predict it. Manager 12, despite being recently joined, leads a high-output team and describes exactly the same open, safe environment as the long-tenured high-output managers. Conversely, Managers 5 and 15, despite being long-tenured, lead low-output teams and describe the same transactional, efficiency-focused style as the recently joined low-output managers. This strengthens the case considerably: it's not tenure itself driving output, it's the work environment and management style that makes the difference.
- Second, several long-tenured managers speculate directly about why that style differs: they point to a shift in management style: older managers extending trust and autonomy, newer managers being more hands-on, metrics-driven, and in a few cases less personally connected to the team and the product. Worth noting, this isn't unanimous.
- Together, this doesn't just confirm the organizational data's tenure/output correlation, it suggests a candidate mechanism: not tenure itself, but the work environment and leadership style associated with it.

Case

Jess and her team gather to go through the interview findings. For a moment, after laying it all out, nobody says anything. Then one of the team members breaks the silence. "Oh my god," she says. "They're hiring the wrong managers!" Jess asks her to explain. Looking back over the interview summaries, the team member points out a pattern: over the past two years, most of the managers Silverback has brought in fit a similar profile — transactional, goal-oriented, focused squarely on performance and efficiency. That style, she argues, is precisely what discourages people from bringing new ideas forward in the first place. If that's what's been getting hired into leadership roles, it could explain why the number of logged changes and patents has been sliding.

Jess nods slowly. "That's a plausible logic model," she says. "But it's still just a hypothesis." She pauses, thinking it through further. "It might not stop at the managers. If Silverback's

recruitment has shifted toward valuing performance and output over creative or innovative behavior, that same shift could apply just as easily to the people being hired." This isn't something the team can answer from the evidence currently available. Jess picks up her phone and calls her contact in Silverback's HR department, putting the call on speaker so the rest of the team can listen in.

When Jess asks, the HR contact explains that recruitment and selection at Silverback has actually been run by Sony ever since the acquisition — but the underlying selection criteria themselves haven't changed. They're still based on a framework Ella herself set up in the earliest days of the studio. Ella's original framework, the HR contact explains, distinguished two broad types of employees. The first she called the "creators" — the lead programmers, systems designers, technical artists, and graphics engineers. These are the people whose job is fundamentally to be creative: to come up with new ideas, spot smart improvements, push the technology forward. Ella considered them the core assets of the studio, and built selection criteria around it accordingly — a proactive personality, strong learning goal orientation, high innovative self-efficacy, openness to new experience, that kind of profile. The second group she called the "worker bees" — the people whose job is to reliably execute: junior and mid-level programmers, QA testers, build engineers, localization specialists, production coordinators. For these roles, the selection criteria have always looked quite different: conscientiousness, attention to detail, reliability under deadline pressure, a strong ability to follow process and instructions closely.

None of that, the HR contact says, has changed. However, what has changed is how managers are selected. About two years ago, Sony concluded that given Silverback's explosive growth, the studio needed what they called "true" managers — preferably MBA-trained, people who could steer teams on output, performance, and getting things done, with a more transactional leadership style than the typical 'people managers' the studio had relied on before. Most of these managers have extensive experience in sectors such as finance, insurance, logistics, and manufacturing; only a few have experience in the gaming industry. That shift in the management hiring criteria was rolled out roughly two years ago.

Jess ends the call and sets her phone down. For a moment, the room is quiet again — but this silence has a different quality than the last one. The team member who'd made the "wrong managers" comment is grinning. She looks, more than anything, like someone who's just been proven right. Before anyone can make a comment, Marian, the team member who has conducted the scientific literature review cuts in. "Before you all roll over each other congratulating yourselves," she says, "maybe take a look at this first. It’s a summary of the main findings of my review of studies on innovative behavior. In particular, pay attention to the construct referred to as psychological safety."

Question

Download the summary of the scientific literature review on innovative behavior <<here>>. What do you consider the most relevant findings for the case?

Answer/feedback

For the Silverback case specifically, three findings stand out:

1. **Psychological safety** is the predictor (r = .30 individual, r = .44 team-level) that matches directly with the manager interviews: high-output teams described exactly this construct ("no bad suggestions," no punishment for "crazy" ideas), low-output teams described the opposite. The suggested mechanism — psychological safety as the pathway through which leadership style translates into innovative behavior — maps precisely onto the "wrong managers" hypothesis: it's not about the managers themselves, but about whether they create (or fail to create) psychological safety within their team.
2. **Leadership style** is the second predictor relevant to the case. The literature doesn't simply indicate that 'better' leaders produce more innovation — it points to a specific 'style': only leaders who promote autonomy, demonstrate fairness, and foster psychological safety create the conditions that drive innovative behavior.
3. **Employee characteristics** such as proactive personality and learning goal orientation are identified as the most robust individual-level predictors, outperforming psychological safety and leadership style in some comparisons. However, according to Jess's HR contact, this factor has already been accounted for: Silverback's "creators" have always been selected on exactly these traits, and that selection criteria hasn't changed. This effectively rules out individual-level hiring as the explanation for the decline.

Case

When Marian walks the team through the review and lays out the main findings, there's a moment of genuine disbelief around the table. Rarely does evidence line up this cleanly — three separate sources, gathered independently, all pointing in the same direction. "I want to be clear about what we actually have here," Marian says. "The manager interviews, the organizational data, and now the scientific literature all support the 'wrong managers' hypothesis. That's three strong, credible sources reinforcing each other. But the proof of the pudding is going to be the stakeholder survey — what the creators themselves report."

To make that test meaningful, Marian explains, the survey needs to do more than ask people what they think is going wrong. It needs to actually measure the two constructs the other evidence has pointed to: psychological safety and leadership style.

Psychological safety, she says, is the easier of the two — seven items adapted from Edmondson's (1999) original scale will do the job, it's well-established and validated. Leadership style is trickier. There are dozens of scales claiming to measure it, Marian explains, but most are built around one or two particular styles and their validity and reliability vary considerably. The most researched and best-validated option is the Multifactor Leadership Questionnaire (MLQ, Bass, 1995), but even that has a limitation: it has 36 items and measures transformational, transactional, and passive/avoidant leadership. Marian's proposal is pragmatic — take the MLQ as a foundation, reduce it to a handful of core items that differentiate transactional from transformational leadership, and add a few additional questions specifically targeting empowering leadership, since that's the dimension the case actually needs to test.

Jess agrees, and the team gets to work on the questionnaire. Three days later, it's been pilot tested and signed off. From there, they follow the procedure agreed on earlier: the survey goes out through SurveyMonkey to 915 lead programmers, systems designers, technical artists, and rendering engineers, drawn from both high-output and low-output teams, and from teams under both long-tenured and recently joined managers. Within a week, 693 responses come back — a response rate of roughly 75%.

Question

1. *Open the raw survey data <<here>> and skim it — get a feel for the structure and scale of the dataset before reading the analysis.*
2. *Read the summary of findings <<here>>.*
3. *What are the main findings? For each one, briefly state how it relates to (confirms, adds to, or complicates) what emerged from the other three evidence sources*

Answer/feedback

- **Psychological safety and leadership style predict team output; manager tenure largely doesn't.** This finding confirms the "wrong managers" hypothesis from the interviews, but also adds focus: it's leadership style, not tenure itself, that matters.
- **Transformational and empowering leadership are the strongest predictors, transactional leadership has a negative impact.** This finding matches the contrast

that emerged from the interviews (trust/autonomy vs. efficiency/control) and the scientific literature's findings on leadership working through autonomy and safety.

- **Psychological safety partially mediates leadership's effect.** This finding directly confirms the scientific literature's finding that psychological safety is a key mechanism linking leadership to innovation — but also shows it's not the whole story, consistent with the literature flagging other pathways too.

Case

After Jess and her team have gone through the survey findings, checking them against everything gathered so far, their conclusion is unambiguous: this fourth source closely aligns with the other three - practitioners, organizational data, scientific literature, and now stakeholders themselves — all four point in the same direction. Taken together, the aggregated evidence for the "wrong managers" theory is about as strong as it gets. The likelihood that something else is driving the decline in changes and patents is small enough that Jess is comfortable treating this as the root cause, not just a leading hypothesis. Which leaves the team facing the final question: now what?

Step 6. Applying the evidence

Case

Jess and Marian meet with Bart, Joel, and Ella to walk them through the findings from all four sources. When they finish, it is quiet for a few moments, then Ella breaks the silence. "We're hiring the wrong managers", she says. "That's exactly the conclusion we came to," Jess says.

Ella's expression shifts from disbelief to something closer to anger. "How did this happen? When Sony took over recruitment and selection, we agreed they'd use the criteria I developed."
It takes a few minutes of back and forth to piece together what actually happened. For the employees — the criteria differentiating between "creators" and "worker bees" are still there, untouched, exactly as Ella had originally designed it. But somewhere along the way, Sony has introduced its own criteria for management hires: people who steer teams on output, performance, and getting things done, with a distinctly transactional leadership style.

Joel sighs and rubs his face. "So what now? We can't exactly fire every manager who fits that profile and replace them with... what, exactly?" "Right," says Jess, "but even with the right criteria in place, a meaningful proportion of new hires will still turn out to be the wrong fit, wrong leadership style. Selection helps, but it's not perfect." She pauses. "But hiring isn't the only lever you have here. There are other options."

Marian picks up where Jess left off. "What Jess means is training and professional development," she says. "It's the standard response when new hires don't quite meet expectations — in most cases you can’t simply replace them, so you develop them." She glances around the table. "Now, to be clear, you're not going to turn hard boiled, transactional, output-driven managers into transformational people-leaders. But you can teach them why delegating authority matters, why building a psychologically safe environment matters for the kind of work this studio depends on. And once they understand why, you can train them in how to actually do it." "That conclusion," she explains, "is what sent us back to the scientific literature. This time we focused on the question: what actually drives psychological safety in a team? If training is going to be the intervention, we need to know exactly what the content of that training should be, and what management skills it should actually build."

CAT Assignment JG 23

1. *Search for empirical studies on the antecedents (predictors/drivers) of psychological safety in teams and/or organizations*

2. *List the database(s) and search terms you used, and report the results of your search queries.*

3. *Select the three most relevant and methodologically strongest studies (preferably meta-analyses and or longitudinal studies) and summarize their key findings in a table, including: year of publication, population, sample size, study design, main findings, effect sizes, and limitations.*

4. *Provide an overview of the most relevant findings, including contextual and moderating factors.*

5. *Based on these findings, what elements would you include in a training for Silverback's managers, aimed at helping them build a more psychologically safe team environment?*