STEALTH NOTES // ISSUE 02 // JUNE 2026

The IRR math, and why the 3.5× premium isn't what you think.

By Ali Sina · Founder, STEALTH

READ TIME: ~10 minutes

COMPANION TO: The State of the Venture Studio Economy 2026 + Fund I LP Deck

WORD COUNT: ~1,600

> *“The 3.5× premium is real. The story underneath it is more interesting than the headline.”*

In Issue 01 I cited a number. Studio-built companies generate roughly a 3.5× IRR premium over traditional seed-funded peers, according to the GSSN / Mind the Bridge multi-studio cohort study. I cited it the way it's usually cited — as a single, clean, headline-friendly fact.

Several LPs read Issue 01 and wrote me back the same question, in different ways. *Is the 3.5× actually real, or is it selection effects?*

This issue is the longer answer. It is more interesting than the short one.

What the number actually says

The 3.5× comes from a 2022 study by the Global Startup Studio Network in partnership with Mind the Bridge. The study analyzed 88 venture studios, 2,341 portfolio companies, and compared them on a like-for-like basis against PitchBook's seed-stage cohort data for the same vintages (2013–2018).

The methodology I used in our Annual Report inherits this study's numbers because the alternative — collecting fresh primary data across 88 studios — would have delayed publication by twelve months.

Two facts about the study that the headline doesn't tell you.

First: the 3.5× is an *average* across the participating studios, not a median. The distribution is wide. Top-quartile studios run closer to 5×. Bottom-quartile studios are closer to 1.5× — meaningful outperformance, but not the headline.

Second: the studios in the study self-selected into the GSSN's reporting framework. There is no reason to assume the studios that did not participate had performance identical to those that did. Most likely the participating cohort skews stronger.

This is the right place to admit something every researcher in venture analysis admits in private: *almost every IRR comparison in this asset class has selection effects.* Traditional venture benchmarks have the same problem in the opposite direction — the funds that fail quietly do not report their final numbers. The 3.5× should be read as directional, not precise. It almost certainly does not represent the marginal new studio launched in 2026.

What I think the real number is

If you correct for selection bias and reset against the broader category — including the bottom-quartile studios that did not participate in the original cohort — I think the honest premium is closer to **2× to 2.5× over traditional seed**, not 3.5×.

I am happy to defend this in writing.

Two-point-something is still a meaningful outperformance. It is also a number that survives the kind of due diligence a serious institutional LP will run on it. A 3.5× claim that collapses to 2.2× when an allocator stress-tests it is worse than a 2.2× claim that holds up. The first one ends the conversation; the second one starts it.

I would rather pitch on 2.2× and have it hold than pitch on 3.5× and have it argued away.

Why studios still outperform — the part that does survive scrutiny

The IRR premium, however measured, comes from three structural mechanisms, all of which survive selection-effect correction:

1. Lower formation-stage failure. Studios prevent the most common early-stage failure modes — premature scaling, wrong-hire founder/market fit, mismanaged runway. The data is consistent across multiple cohort studies. Studio-built companies have a roughly 30% lower failure rate by month 18. This is the largest single source of seed-stage capital destruction, and the studio's operational scaffolding addresses it directly.

2. Faster milestone velocity. Studios compress the time from formation to first material revenue by approximately 40%. The mechanism is observable: pre-built distribution, pre-existing hiring funnels, pre-validated supplier relationships. The first six months of a studio-built company is operationally different from the first six months of a typical seed-funded startup. This is not selection bias; this is structural.

3. Higher Series A clearance. Studio-built companies clear Series A at a roughly 25% higher rate than seed-funded peers at similar vintage and sector. This effect compounds the first two — if you survive month 18 and you hit revenue 40% faster, you are simply more likely to be a credible Series A target.

The three mechanisms together generate the premium. They also explain why the premium will not stay at 3.5× as the category matures. As more studios launch, the operational practices that produce these advantages will diffuse across the broader venture ecosystem. The premium will narrow. The studio model will not stop working — it will just stop being unusually well-priced relative to the work it does.

What this means for an LP modeling Fund I

If you are evaluating a studio fund commitment, do not anchor on the 3.5× number. Anchor on the three mechanisms. Ask the GP:

Failure rate at month 18. What is yours? How is it measured? Across what vintage? If they cannot answer this without consulting the operating team, that is a signal.

Time to first material revenue. Defined how? Median or mean? Across all formed companies or only the ones still operating? If the answer is the latter, you are looking at survivor bias.

Series A clearance rate. What percentage of formed companies have cleared Series A inside their cohort vintage's typical window? This is a metric the GP either owns by heart or does not. It is the cleanest single test of operational quality.

If you get clean answers to those three, the premium is real for that studio regardless of what the GSSN average looks like. If you do not, the 3.5× is just a marketing number.

What we measure at STEALTH

For the record — the data I am willing to put in writing about our own portfolio, knowing this is going to be read by people who care about the difference between marketing and underwriting:

Our 18-month failure rate, measured across the 2020–2023 vintages, is approximately **32%**. The baseline for seed-funded peers in the same vintages, per PitchBook, is approximately **55%**. We are roughly 23 percentage points better, or about a 42% relative reduction.

Our median time from formation to first $25K MRR, measured across the same vintages, is approximately **11 months**. The baseline for seed-funded peers is approximately **17 months**. We are roughly 35% faster.

Our Series A clearance rate within 36 months of formation, measured across the 2020–2022 vintages (where 36 months has elapsed) is approximately **24%**. The baseline is approximately **17%**. We are roughly 40% higher in relative terms.

These numbers are not selection-effect-corrected. They are the rawest version of our own internal performance, dated June 2026. They will be refreshed quarterly in the Distributed Studio Index as we publish it.

What I will not pretend

I will not pretend these numbers will hold across our next 700 companies. The portfolio expansion from 700 to 1,400 will probably regress some of these metrics toward the broader-category mean. The mechanisms that produce outperformance are durable; the magnitude is not guaranteed at scale.

I will also not pretend that any single GP's claim about their own portfolio is independently verifiable. You are taking my word for it that our 18-month failure rate is 32%. The peer review I would actually want for this number

is participation in the Distributed Studio Index (DSI) — an external benchmarking framework with cross-studio audit. We are launching the DSI in Q4 2026. By Q4 2027, you will be able to compare our self-reported numbers against the index's aggregate, and against the other distributed studios that participate.

Until then, the best I can give you is the methodology, the source studies, and the willingness to be specific in writing.

What I would tell you if you were modeling this

Use 2.2× as your premium assumption. Stress-test against 1.8× — i.e., assume our mechanisms produce a smaller advantage than we currently measure. Confirm that your base case clears your hurdle at 1.8×. If it does, the Fund I commitment is defensible regardless of how the headline numbers settle.

If 2.2× is enough — and for most institutional allocators with a 10-year horizon and a defined alternatives bucket, it is — then the headline premium is the wrong thing to spend time arguing about.

The interesting question is not how big the premium is.

The interesting question is whether the mechanisms that produce it are durable.

I think they are.

The DSI will tell us how right I am.

— Ali Sina
Founder, STEALTH
ali@stealth1000.com · stealth1000.com

> **Next issue:** *the operating math behind compounding. How company 300 actually gets cheaper to build than company 30, in dollars and weeks. Subscribe at stealth1000.com/research.*