A FOUNDERS' BRIEF

Failure Intelligence

What 2,715 startup deaths
actually tell us about
commercial validation

2,715
documented
startup deaths

$660B
capital traceable
across the records

58%
of traceable capital
died at Stage 3

A founders' brief, built from 2,715 documented startup deaths across six public sources.
Read it before you build, before you raise, before you hire.

01

Before you read this

Founders are now able to build faster than they can think. A landing page, a prototype, a demo deck and a pitch can all be assembled in a weekend. That is genuinely useful. It also means the part of a business that breaks first, the commercial logic, is now the part most likely to be skipped.

The 2,715 records used here are public post-mortems gathered from six sources: Loot Drop, the Kaggle CB Insights failure dataset, the CB Insights failure database, and Failory's startup, Google and Amazon cemeteries. Each record was de-duplicated and tagged with a GTM Right stage where the documented cause of death allowed a clean mapping. Funding figures, country, sector, founding year and year of death were retained where present.

What this report does and does not claim

The dataset is useful for pattern recognition. It is not useful as proof that any specific intervention would have prevented any specific failure. Three honest limits matter, and they shape every percentage in this document:

- **The mapping is mechanical, not interpretive.** If a post-mortem said "Competition," the record was tagged Stage 5. If it said "Unit Economics," Stage 3. The mapping is consistent and reproducible. It also inherits whatever oversimplification sits inside the original source.
- **956 of the 2,715 records (35%) sit in Unknown.** They had no cause information clean enough to assign a stage. Almost all of these come from CB Insights and Failory entries that captured the company but not the diagnostic narrative.
- **This is a graveyard.** We do not see the survivors. The report can describe how companies died. It cannot describe how the same diagnostics performed inside companies that lived.

Read the rest as a map of where the cliffs are, not as a guarantee that the right diagnostic keeps anyone off them.

02

The GTM Right framework

Every commercial system from idea to profit passes through six gates. Each gate is a different question. A startup that fails commercially has almost always failed at one of these six, even when the post-mortem labels it something else. The framework below is the map this report uses to read the dataset.

The stages are sequential in importance, not in time. A founder with a finished product can still have an unanswered Stage 1 question. A company at Series C can still die from a Stage 2 problem its earlier funding rounds papered over.

Stage	Name	The question it tests
1	**True Validation**	Is there a real, repeatable pain a buyer will describe in their own words?
2	**Buyer Precision**	Can you name the first reachable buyer specifically enough to find them this week?
3	**Revenue Funnel**	Can one customer be made profitable without invoking future scale?
4	**Buyer + Offer**	What job is this product hired to do, and what does it replace?
5	**Market + Funnel**	What underserved wedge can you hold against incumbents and distribution gravity?
6	**Commercial Decision**	Could one regulatory, legal or trust issue block launch, sales or adoption?

Why these six and in this order. Stages 1 and 2 establish that there is a buyer and you know who they are. Stage 3 establishes that serving them is a profitable act. Stage 4 establishes the offer's fit against the actual job being done. Stage 5 establishes a defensible position in a market with other players. Stage 6 establishes that no external decision-maker, a regulator, a court, a counterparty, can shut the system down. Skip any one of these and the failure mode that follows is predictable.

The rest of the report reads the dataset against this map. Where the data is clean, the picture sharpens. Where the data is muddy, the muddiness is itself a finding, and Section 04 deals with that directly.

03

Where the money actually dies

Of the $660 billion in known capital across the 1,759 records with funding data, $384 billion (58%) was committed to companies that died at **Stage 3, the Revenue Funnel stage**. That is more than five times the next most expensive stage.

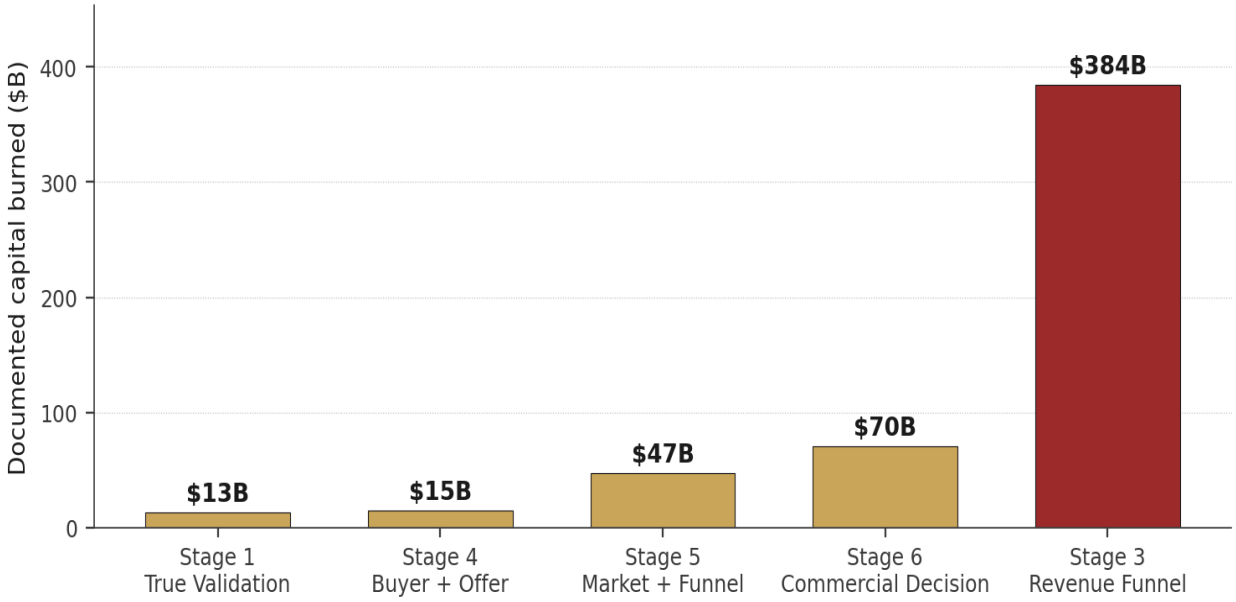

Documented capital totals across all records with funding data, grouped by the GTM Right stage where the company died. Excludes 956 Unknown-stage records and 492 records with no funding data.

Stage 3 is the question of whether one customer can be made profitable. Whether the cost of attention, conversion, delivery and retention adds up to less than the price the buyer will pay. It is what dies inside WeWork ($22B raised, unit economics never closed), Northvolt ($15B, manufacturing cost structure), Byju's ($6B, paid acquisition far ahead of retention), and dozens of others.

The asymmetry that founders should care about

Failures cost different amounts of money depending on where they happen. The median funding at death tells the story:

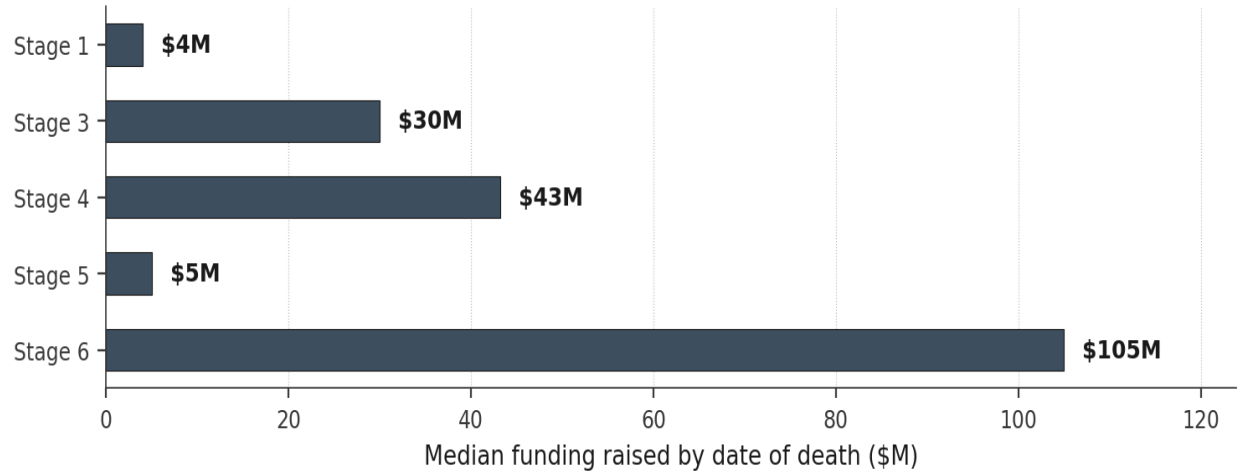

Median funding raised by date of death, by GTM Right stage. Computed across records with positive funding figures only.

The implication is concrete. Validation failures (Stage 1) cost a median of $4M. Stage 5 competitive deaths cost a median of $5M. Stage 3 deaths cost a median of $30M, and the worst outliers go past $20B. The hour you spend modelling unit economics before raising returns more capital, more time and more reputation than the hour you spend doing almost anything else upstream of build.

FOUNDER TAKEAWAY

Cheap failure usually happens at validation. Expensive failure usually happens when the funnel economics were assumed rather than proven.

04

How the cause-of-death labels lie

A founder writing a post-mortem is choosing one sentence to put on a headstone. The sentence almost always points at a downstream symptom rather than the upstream cause. Two of these mislabels are systematic enough to bend the whole dataset, and a careful reader needs to see them before reading any of the percentages that follow.

Mislabel one: “Competition” hides Stage 1 and Stage 4

52% of all clean-cause records say the company died because of competition. This is the most common, the most quoted, and the most useless cause of death.

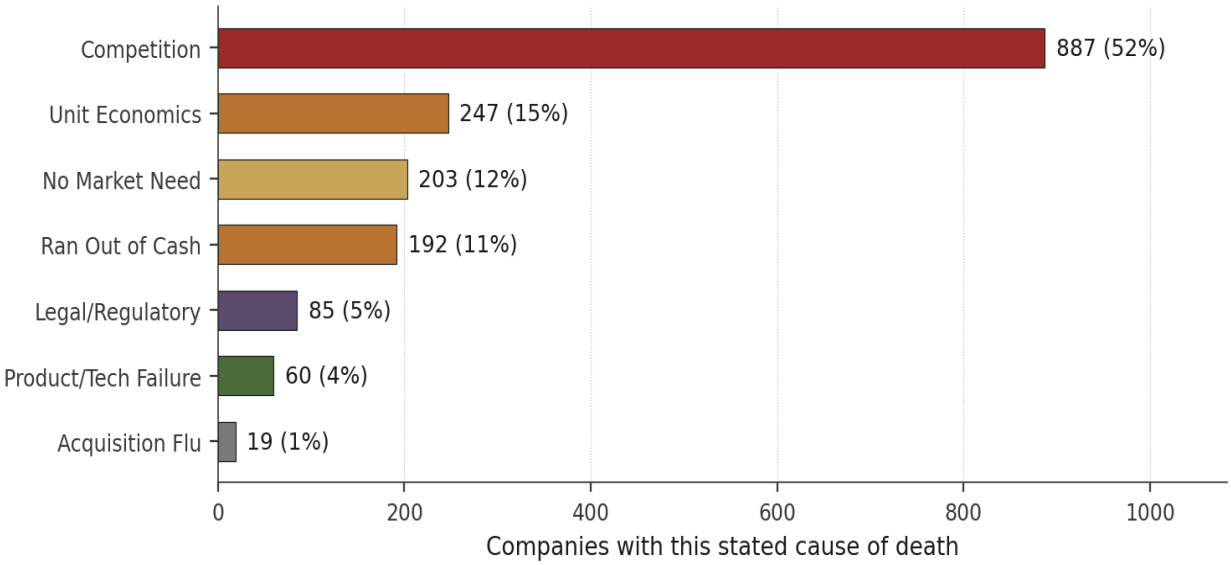

Documented causes of death across the 1,705 records with a clean primary cause. Percentages are of known causes only.

“Competition” is what founders write when they cannot put a sharper sentence on the headstone. Almost every Stage 5 competitive death is downstream of one of three earlier failures:

- The pain was not sharp enough to make a buyer switch (Stage 1)
- The offer was not structured around a job the buyer was willing to pay for (Stage 4)
- The wedge into the market was never identified (the actual Stage 5 question)

Reading “we were beaten by competitors” as the diagnosis is reading the symptom. The diagnostic question, the one that would have changed the outcome, is the one that should have been answered before the build:

“What did we never establish that would have made competition irrelevant?”

04 (continued)

How the cause-of-death labels lie

Mislabel two: Stage 2 is the silent graveyard

Only one record in 2,715 was tagged Stage 2 (Buyer Precision). The honest reading of that number is not that Buyer Precision failures are rare. It is that no founder writes “I did not know who I was selling to” on the headstone. That sentence is too uncomfortable to say out loud, so it gets relabelled as something else.

Stage 2 failures hide inside two more flattering labels:

- **“No Market Need” (Stage 1).** The market often did exist. The team simply could not point at a specific buyer they could reach in a specific way. The pain was real; the buyer was a fog.
- **“Bad Marketing” or “Bad Distribution” (Stage 4 / Stage 5).** The marketing was not bad. It was aimed at no one in particular, because no one in particular had been named. Channels did not fail; targeting never existed.

This is why Stage 2 is the silent graveyard of the dataset. Its real population is almost certainly large and is distributed across the Stage 1 and Stage 4 columns of every chart in this report. A founder reading those columns should assume that some unknown share of them is a Buyer Precision failure wearing a Stage 1 or Stage 4 mask.

The diagnostic question that would have surfaced it earlier is unsentimental:

> **“Can I name, today, the first ten reachable buyers, the trigger event that puts them in-market, and the channel that finds them this week?”**

How to read the rest of this report

Both mislabels work in the same direction. They protect the founder from naming the question they did not answer. “Ran Out of Cash” is rarely the cause; it is the consequence of unit economics that never closed. “No Market Need” is sometimes accurate, and sometimes a polite description of a buyer who was never identified. When you read a startup post-mortem, treat the stated cause as the starting question, not the answer.

05

Failure shifts as capital scales

Different funding levels die different deaths. Two-thirds of failures under $10M die at Stage 5, where the company built something, raised early money, and never broke through against incumbents or distribution gravity. By the time capital crosses $50M, the dominant failure stage shifts decisively to Stage 3. Half of all failures above $100M happen there.

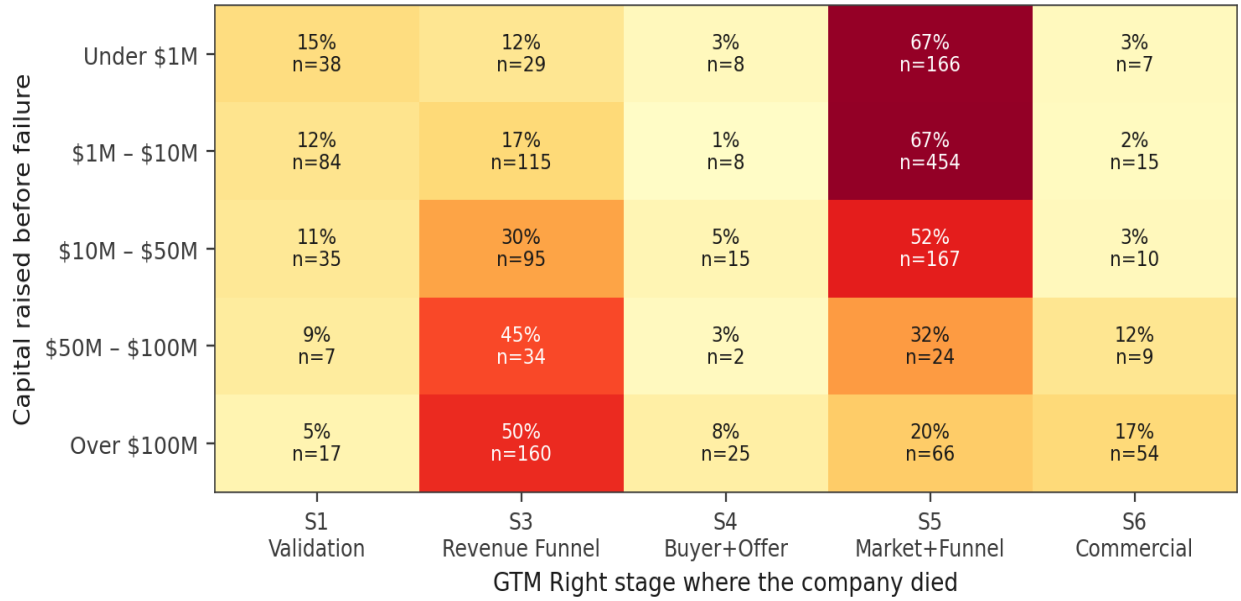

Each row reads as 100% of failures within that funding band. Cell values show the share of failures in that band that died at the named stage, with raw counts. Excludes Unknown-stage and no-funding records.

This is not because validation matters less for well-funded companies. It is because validation failures are caught earlier, cheaper, and quieter. The $100M company that dies has typically passed Stage 1 already, and is dying because the unit economics it promised investors never materialised at scale.

What this means for the founder

Stress-test the question most likely to kill you at the funding level you are actually trying to reach. A founder raising under $10M should rehearse Stage 5: the wedge, the distribution path, the reason the first hundred buyers choose you over an incumbent. A founder raising past $50M, where the survival question becomes “do the unit economics actually work?”, should already have answered Stage 3 with real numbers, not assumed margins.

The 17% Stage 6 share at $100M+ is also worth noting. Late-stage failures from regulatory, legal or acquisition-flu issues account for one in six deaths above $100M. They are absent at low funding levels, not because they cannot happen, but because companies that small rarely live long enough to encounter them.

06

Each sector dies a different way

Stage averages flatten the picture. The sector view is sharper, and the pattern in each one tells a founder which question matters most for their kind of business.

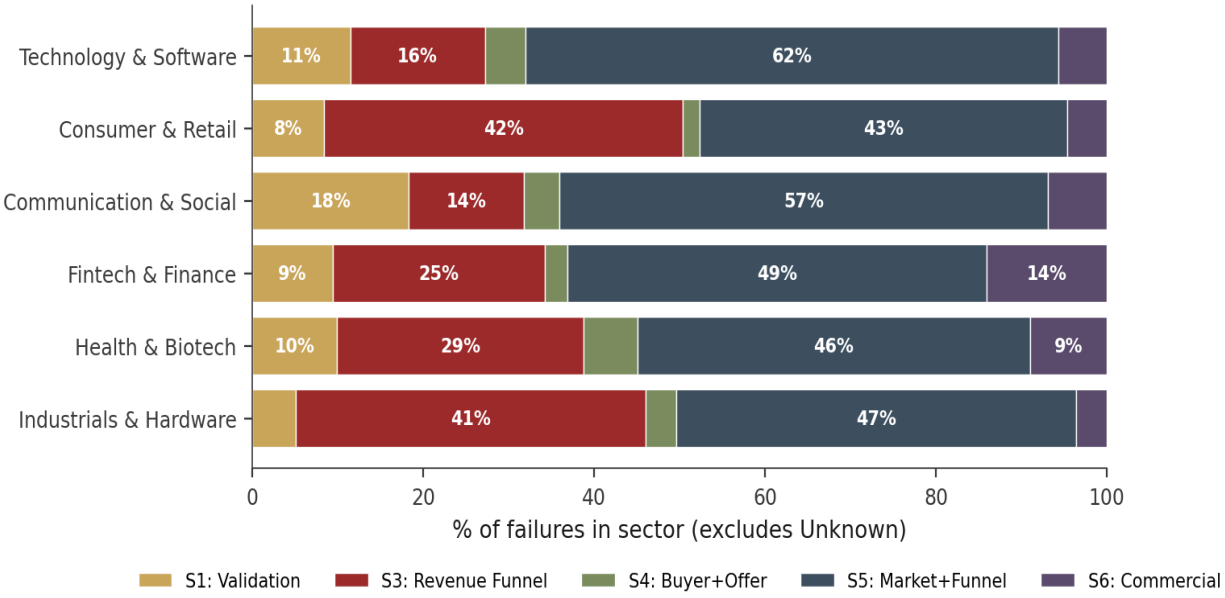

Distribution of failure stage within each of the six largest sectors. Bars are normalised to 100% of mapped failures in that sector; Unknown-stage records are excluded.

Sector	Records	Where it tends to die	What to test hardest
Technology & Software	637	Stage 5 (62%)	Wedge, defensibility, distribution
Consumer & Retail	456	Stage 3 + Stage 5 split	Unit economics first, then channel
Communication & Social	441	Stage 5 (57%); Stage 1 (18%)	Network formation and real pain
Fintech & Finance	177	Stage 5 (49%), Stage 3 (25%), Stage 6 (14%)	CAC, regulatory pathway
Health & Biotech	162	Stage 5 (46%), Stage 3 (29%)	Reimbursement, clinical adoption
Industrials & Hardware	137	Stage 3 (41%), Stage 5 (47%)	Manufacturing cost structure

Two patterns are worth lingering on. Communication and Social is the only sector where Stage 1 (no validated pain) breaks 18%, which fits the experience of network-effect markets: many of these companies died because the pain they imagined was never sharp enough to displace the network the buyer was already inside. Fintech is the only sector where Stage 6 (regulatory, compliance, fraud) breaks 10%. If you are building in finance, that line is not optional reading.

07

Where you build shapes how you fail

Most failure reports treat geography as a footnote. The dataset suggests it should not be. Each major venture ecosystem has a distinctive failure signature, and a founder rehearsing the wrong question for their geography is doing diligent work on a problem they are not actually going to face.

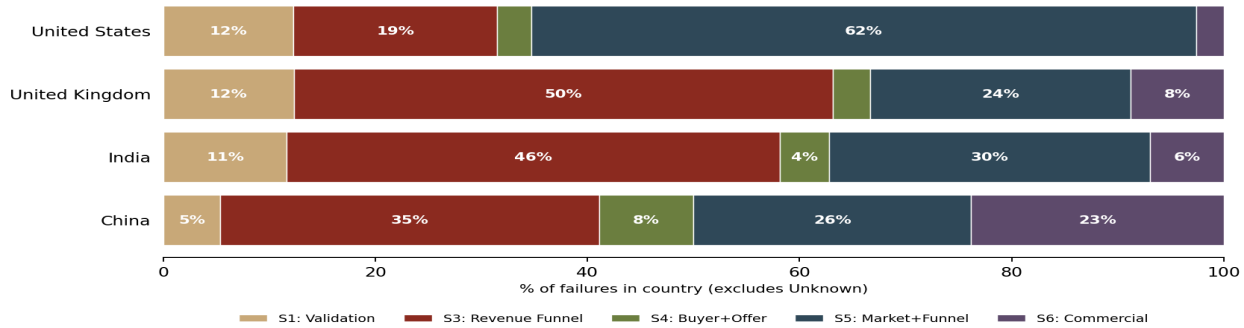

Distribution of failure stage within each of the four largest geographies. Bars are normalised to 100% of mapped failures in that country; Unknown-stage records excluded. Sample sizes: US n=1,225, China n=168, UK n=57, India n=43.

United States: a market-and-funnel graveyard

62% of mapped US failures died at Stage 5. American failure is overwhelmingly competitive. Capital is plentiful, product ships quickly, and the question that kills companies is whether they could hold a defensible wedge in a crowded market. Uber Rush, iRobot, Jawbone, Groupon Now, Boxed, Tellme, Tintri: each entered a contested lane with a better mousetrap and discovered that better was not the variable that mattered.

China: a regulatory-and-capital graveyard

23% of mapped Chinese failures died at Stage 6, roughly eight times the US rate. Another 35% died at Stage 3. The regulatory column is the China-specific signature: the 2021 education-policy reform alone wiped out Yuanfudao, Zuoyebang, Dali Education, VIPKid and Zhangmen, with billions in combined funding. The Stage 3 column is the second story: capital-intensive consumer plays, particularly in EV and auto (LeEco $6B, WM Motor $5.8B, Qoros $3B, Singulato $2.4B), where unit economics never closed at the scale capital required.

India and the United Kingdom: revenue-funnel graveyards

50% of mapped UK failures and 46% of Indian failures died at Stage 3. Both ecosystems show the same dominant pattern: a buyer existed, a product shipped, and the unit economics never closed. In India, this is the Byju's and WhiteHat Jr story: paid acquisition that ran ahead of retention. In the UK, it is a smaller-scale version of the same problem, often visible at the late seed and Series A stages where companies raised once at strong valuations and could not raise again on the underlying numbers.

The diagnostic question that matters depends on where you are building. A US founder rehearsing Stage 5 wedge and distribution questions early is doing the right work. A Chinese founder weighting Stage 6 (regulatory, licensing, policy) is doing the right work. An Indian or UK founder who has not yet answered Stage 3 (one customer, profitable, today, on paper) is rehearsing the wrong question.

08

Six named patterns and their cost

Aggregate numbers fade. Named cases stick. Below are six recurring patterns from the dataset, each with the companies that lost the most money to them, and the question the founder should have answered earlier.

STAGE 1

The Painless Product

A product the founder believed should exist, without enough evidence that buyers were already trying to solve the problem.

Named cases in the dataset: Hopin ($1.07B, 2024). Udacity (~$1B, 2024). Better Place ($850M, 2014). Coinbase NFT ($800M, 2024). Pico ($1.5B). Wolfspeed ($2.0B).

Founder check. Can you cite ten real instances of a buyer describing this pain in their own words?

STAGE 2

The Invisible Buyer

A real market existed, but the team never named the specific buyer they were selling to. Spend went out to no one in particular and the post-mortem was filed under “No Market Need” or “Bad Marketing.”

Named cases in the dataset: Essential Products ($330M, premium Android with no defined buyer who was not already on iOS). Doppler Labs ($50M, smart earbuds aimed at “audiophiles and tech enthusiasts”). Color Labs ($41M, Sequoia-backed photo app with no specified user). Mode Media ($230M, lifestyle network whose advertiser-buyer was never named). Beepi ($150M, P2P used cars unsure whether buyers, sellers or both were the customer).

Founder check. Can you name the first ten reachable buyers, the trigger event that puts them in-market, and the channel that finds them this week?

STAGE 3

The Broken Funnel

Unit economics that only worked at scale, where the scale itself required capital that required the economics to work first.

Named cases in the dataset: WeWork ($22B). Northvolt ($15B). Byju’s ($6B). LeEco ($6B). Convoy. Quibi.

Founder check. Can you describe how one customer becomes profitable without invoking future scale?

STAGE 4

The Feature Without a Job

A product with users who valued it but would not pay, or pay enough, because the offer did not match a job the buyer was willing to fund.

Named cases in the dataset: Theranos ($700M, the technology never delivered the offered diagnostic). Fire Phone (3D features that added cost without adding buyer value). Google Glass.

Founder check. What is this product being hired to do, and what does the buyer stop using or replace once it works?

08 (continued)

Six named patterns and their cost

STAGE 5

The Crowded Lane

Entered a contested market without a clear underserved segment to anchor on. Being better than incumbents was not enough to make a buyer switch.

Named cases in the dataset: Quirky (~$185M, consumer invention platform). Knewton ($182M, adaptive learning). Zhenai ($200M, dating). Faraday Future ($3.5B). Uber Rush.

Founder check. Can you name the gap competitors are not solving, and the precise reason you can win it?

STAGE 6

The Regulatory Blindspot

A business that was viable on paper and uneconomic once compliance, licensing, or financial-services obligations were priced in. Or where a single legal event changed the operating environment overnight.

Named cases in the dataset: Wirecard ($28B, payments fraud). Ezubao ($7.6B, P2P lending crackdown). Yuanfudao ($4.1B), Zuoyebang ($2.9B), both hit by China's 2021 education-policy reform.

Founder check. Could one regulatory, legal or trust issue block launch, sales or adoption?

What ties the six patterns together

In every named case, the failure was not the moment the company shut down. It was the moment, often years earlier, that a commercial assumption stayed untested while the team kept building around it. The six patterns are not about prediction. They are about which assumption, in your specific business, is currently doing the most work without enough evidence underneath it.

A founder using this list well does not score themselves green on all six. They identify the one or two patterns most likely to apply to their situation, and they put their next two weeks of work into producing evidence that resolves them.

09

What the data does not tell us

A credible failure report is honest about its blind spots. Four matter for how the percentages above should be read.

Section 04 already covered the largest one: that Stage 2 (Buyer Precision) is systematically underreported, with most of its real population hidden inside Stage 1 and Stage 4 columns. The four limits below sit alongside it.

Source weight matters.

Loot Drop accounts for 1,675 records, 62% of the total. Its tagging conventions therefore set the pattern. The 442 CB Insights records had no cause field clean enough to map and sit in Unknown. Patterns observed here are partly patterns of how Loot Drop classifies, not only how startups die.

Geographic skew.

54% of records are US companies. China, India and the UK each contribute roughly 3–6%. Patterns from outside major venture ecosystems are under-represented. A founder building in Lagos, São Paulo or Jakarta should apply these patterns directionally and weight local context heavily.

Founding-year bias.

Year-of-death is concentrated 2018–2024. The dataset is mostly a recent-cohort view, with a heavy tail in 2020 (the COVID disruption) and 2022–2024 (the venture correction). Pre-2010 founder behaviour is not well represented.

Survivor blindness.

This is a graveyard. We have no comparison to companies that ran the same diagnostics and survived. A pattern common in failed companies may also be common in successful ones, in which case it is not predictive on its own. The patterns are useful as questions to ask, not as filters to predict.

These limits do not invalidate the patterns. They are why every percentage in this report should be read as directional, and why a sensible founder uses these signals to sharpen questions, not to settle them.

10

A 30-minute pre-build diagnostic

Before the next dollar, the next hire, or the next deck, sit with the idea for half an hour and answer six questions in writing. Not in your head. Written, dated, and re-readable in a month.

Stage	The question	What “yes” looks like
1. True Validation	Can you cite ten real instances of a buyer describing this pain in their own words?	Quotes you can attribute. Not paraphrases, not your own synthesis.
2. Buyer Precision	Can you name the first reachable buyer specifically enough to find them this week?	Company size, role, trigger event, and a list of ten specific people.
3. Revenue Funnel	Can you describe how one customer becomes profitable without invoking future scale?	Cost to acquire, deliver, retain. Less than price. With margin. On paper, today.
4. Buyer + Offer	Can you say what this product is being hired to do, and what it replaces?	A job-to-be-done the buyer pays for today, in some form, with someone.
5. Market + Funnel	Can you name a competitor’s underserved segment, and the wedge you can hold?	A specific group, a specific complaint, and a reason the wedge is yours.
6. Commercial Decision	Could one regulatory, legal or trust issue block launch, sales or adoption?	Mapped, written down, and a named person who has stress-tested the answer.

How to use the result

If three or more answers are softer than they should be, the next step is not building. The next step is evidence. Pick the weakest answer first; that is your bottleneck. Spend two weeks producing the artefact that would let you answer it harder. A list of ten quotes. A funnel model with real costs. A named buyer segment with a list of ten people.

A startup that runs this loop honestly will fail less often, and when it does fail, will fail at a stage where failure is cheap. The dataset behind this report does not contain those companies, because they tend not to feature in public post-mortems. They are absorbed into the survivors.

BEFORE YOU BUILD, RAISE, HIRE OR SPEND

Want to test your idea against these six failure patterns? Run the free GTM Right Validation Gate before the next dollar leaves the account.

11

Methodology

Source reconciliation

Six datasets were ingested: Loot Drop (1,675 records), Kaggle CB Insights startup-failures dataset (510), CB Insights failure database (442), Failory startup cemetery (30), Failory Google cemetery (29), and Failory Amazon cemetery (29). Records were de-duplicated by company name and merged. Funding figures were normalised to USD millions where convertible. Country and sector fields were standardised against a 15-sector taxonomy. Year-of-death was retained where present (2,154 of 2,715 records).

GTM Right stage assignment

Each record's documented cause of death was mapped to the closest matching GTM Right stage using a deterministic rule: "No Market Need" → Stage 1, "Unit Economics" or "Ran Out of Cash" → Stage 3, "Product/Tech Failure" → Stage 4, "Competition" → Stage 5, "Legal/Regulatory" or "Acquisition Flu" → Stage 6. Records with no clean cause field were left as Unknown rather than guessed. The rule is consistent and reproducible. It is also a simplification. "Competition" almost certainly hides upstream Stage 1 and Stage 4 failures, and Stage 2 failures hide inside Stage 1 and Stage 4 columns (see Section 04).

Calculations

Capital totals exclude the 956 Unknown-stage records and the 492 records with no funding data. Median funding per stage is computed across records with positive funding only. Stage percentages by funding bucket and sector are row-normalised across the five primary stages, with Unknown excluded from the denominator. All percentages are rounded to the nearest whole number, which is why some rows do not sum to exactly 100%.

What this report deliberately does not claim

It does not claim that running a GTM Right diagnostic would have prevented any specific failure. It claims that the question being asked at each stage is the question that, if answered honestly before the build, would make the failure visible while the cost of changing direction is still small. That is a different and more defensible claim, and it is the only one the data supports.

Sources

Source	Records	Notes
Loot Drop	1,675	Largest structured dataset. Causes pre-tagged.
Kaggle / CB Insights	510	Detailed founder narratives.
CB Insights	442	Geography and funding depth. Causes not tagged.
Failory Cemetery	30	Consumer and social failures.
Failory Google Cemetery	29	Corporate-scale product deaths.
Failory Amazon Cemetery	29	Corporate-scale product deaths.
Master Dataset	**2,715**	**De-duplicated, GTM Right stage-mapped.**

Run the GTM Right Validation Gate before you build, raise, hire or spend.