

# Revenue that holds through a bad year.

By Rahul Jindal

---

Most revenue is concentrated by default. Sales optimization rewards the highest-converting segment. Customer success rewards the highest-LTV cohort. Marketing rewards the highest-ROAS channel. Each function does its job; the aggregate is fragility.

Resilient revenue is a structure, not a slogan. Diversified revenue is what the deck claims. Resilient revenue is what survives a bad year for one of your axes of concentration. The two are confused so often that the confusion has become the dominant failure mode of growth-stage companies entering their first cycle.

The work is not waiting for the bad year. The work is building the structure during the good year, before the bad year tells you which axis you missed.

*“Diversified revenue is what the deck claims. Resilient revenue is what survives a bad year for one of your axes.”*

## Why most revenue is fragile

Three concentration drives, each running independently inside a healthy growth-stage company.

- 1. Sales concentrates.** Reps optimize for closing, not portfolio. Top-10 logo concentration creeps upward quarter over quarter because the rep who landed the whale is the rep who gets promoted. The CRO sees ARR growth on the dashboard. The quietly rising concentration ratio is not on the dashboard, because nobody is running it.
- 2. Customer success concentrates.** CS teams optimize for the customer in the room, which is the highest-LTV customer. The mid-market customer gets less attention, retains worse, and the cohort thins. Over 18 months, the customer base shrinks toward whales. Net retention looks great because the surviving cohort is healthier than

average; the cohort is also smaller than average, and the underlying base is more concentrated.

3. **Marketing concentrates.** Performance marketing rewards channels with measurable ROAS. The channels with measurable ROAS are usually the most concentrated (Google, Meta, LinkedIn). Brand, partnerships, and word-of-mouth get under-funded because their measurement is harder. The acquisition mix homogenizes; the company is increasingly one channel away from a problem.

The aggregate effect: customer concentration up, vertical concentration up, channel concentration up. The board sees three healthy growth charts. The underlying structure looks like a single bet, dressed up as three.

## **The five axes of concentration**

The bad year, when it arrives, comes from one of five axes. Naming them ahead of the bad year, ranking them, and de-risking the most fragile one is the work.

1. **Customer concentration.** Top-10 customers as a percentage of ARR. Above 30%, fragile. Above 50%, single-point-of-failure. The most measured axis, and often not the most fragile.
2. **Vertical concentration.** Customers in one industry. Banking-software companies in 2008. Travel-tech in 2020. Crypto-services in 2022. AI-application companies whose customers are all VC-funded startups in 2026. The vertical bad year is faster and deeper than the company expects, because every customer in the same vertical hits the same problem at roughly the same time.
3. **Channel concentration.** Revenue through one acquisition channel. Brands dependent on Meta ads when iOS 14.5 ATT shipped. SaaS companies dependent on Google search as AI search reshapes intent. Marketplaces dependent on a single platform. The channel bad year arrives without warning, because platforms change their rules without consultation.
4. **Geography concentration.** Revenue from one country or region. The US-only company that cannot weather a US recession. The China-only company that cannot

weather Chinese regulatory change. The EU-only company that cannot weather a single court ruling.

5. **Product concentration.** Revenue from one SKU, feature, or use case. The single-product company that gets disrupted in its category. The feature-business that gets eaten by the platform feature. The narrow product is fragile not because the product is bad but because every threat lands on the same surface.

*“Most companies de-risk the axis they understand best. The bad year usually comes from the axis they ignored.”*

## Specific moves

1. **Publish the concentration ratios at the leadership cadence.** The CRO reports ARR weekly; the CFO reports gross margin weekly. The five concentration ratios should appear in the same operating cadence, with the same accountability. What gets measured gets defended; what is not measured is what fails first.
2. **Build the second axis while the first is healthy.** Hamilton Helmer's logic on counter-positioning: the time to enter a new vertical is when the current vertical is funding the experiment. The time to fund a new channel is when the current channel is converting cleanly. Companies that wait for the bad year cannot afford the experiment by the time it arrives, because the budget for experiments is the first thing cut when growth slows.
3. **Run the concentration pre-mortem.** Pick the axis you are most concentrated on. Imagine that axis goes flat tomorrow. What does your operating plan look like? If the answer is "we cut costs," you are building fragile. If the answer is "we lean into axes 2 and 3," you have structure. Run this exercise quarterly with the leadership team and the board.
4. **Treat customer success as a portfolio function.** Build CS coverage that retains the mid-market, not just the whales. The mid-market is the insurance policy when a whale leaves. The CS team that optimizes for top-account NPS produces a customer base that looks healthier than it is.

5. **Fund the unmeasurable channels deliberately.** Performance marketing's ROAS is a feature, not the goal. Brand, partnerships, word-of-mouth, and community channels all have worse measurement and better resilience. Allocate against the resilience, not just the measurement. The unmeasurable channels are also the ones competitors cannot replicate quickly.
6. **Distinguish revenue resilience from revenue durability.** Resilient revenue survives the shock; durable revenue compounds over time. They are different design problems. The same company can be resilient and not durable (it survives the bad year but does not grow), or durable and not resilient (it compounds in the good years and breaks in the bad). The leadership team should be explicit about which one they are building, and on what horizon.

## Failure modes

- **Mistaking growth for resilience.** A company growing 80% YoY can have rising concentration ratios on every axis at the same time. The growth masks the structural fragility until both arrive together. The hardest year for a hyper-growth company is its first slow year, because the structural problems compound at exactly the moment the growth stops covering them.
- **Diversifying for the deck, not the field.** Adding a second product that sells into the same customer base diversifies SKUs, not revenue. The customer concentration is unchanged. The board claps for the new product. The structure is no more resilient than before.
- **De-risking the loudest axis instead of the most fragile.** Boards focus on customer concentration because it is the easiest to measure and the easiest to discuss. The bad year usually comes from channel or vertical, which were not on the dashboard. Loud is not the same as fragile.
- **Building resilience after the bad year arrives.** When growth slows and cash gets tight, the budget for de-risking experiments is the first thing cut. The window for building resilience closes precisely when the need for it is most visible. This is why the work has to happen during the good year, even though the good year is when nobody wants to do it.

## Open questions

- What is the right concentration target by axis? "Top-10 below 30% of ARR" is a heuristic, not evidence. Industry-specific ratios are real and under-published. The right target for a vertical SaaS company is different from the right target for a consumer subscription business, and the literature is thin.
- How do compounding engines (from the products page) interact with concentration risk? A network-compounding product is structurally more concentrated than a skill-compounding one; the resilience math has to account for both. The unified framework does not yet exist.
- Is there a structural difference between resilient revenue and durable revenue? My instinct is yes. Resilient survives shock; durable compounds over time. Different design problems, often confused. The taxonomy needs more work.
- How does AI-driven volatility change these ratios? Channel concentration risk in particular looks higher in 2026 than in 2020 because the channels themselves are reshaping faster than companies can rebuild around them.

## The connection

Revenue resilience is the corporate-finance version of the compounding question that runs through the second-sale test (products) and the Adaptive Org (organizations). The work that survives is the work whose value does not depend on the conditions of its first deployment. A product that holds up after the demo, a function that holds up when the founder leaves, a revenue line that holds up when one channel breaks. Same question, three altitudes.

## Related on this site

- Products people pay for the second time: the same compounding question at the product altitude →
- The Margin Thesis: revenue resilience and margin resilience are different problems →
- The Eighth Conversation: operational reliability has the same structural flavor →
- The second-purchase test: same compounding question at the product altitude →

---

Shared privately. Please do not redistribute.