



GROWTH THAT COMPOUNDS

Every growth budget is a gross number

A budget records what will be spent. It does not record what wears out as you spend it, and only the second number says whether the year added anything.

Sinan Isoglu, MBA (Quantic)

Commercial growth leader, lecturer and doctoral researcher

PUBLISHED

1 August 2026

UPDATED

20 August 2026

READING TIME

18 min

WORDS

4,044

<https://isoglu.com/insights/journal/every-growth-budget-is-a-gross-number/>

Management summary	2
Why is maintaining a flat growth budget an active capital commitment?	3
What do foundational studies actually report regarding customer decay rates?	3
Why do customer churn and revenue contraction measure distinct operational assets?	5
Why does commercial management lack external benchmarks for relational equity decay?	6
Why can only internal operating ledgers isolate true baseline customer decay?	7
Why does conventional cohort reporting conceal underlying decay mechanisms?	7
What core econometric baseline survives rigorous retention modeling?	8
Under what specific contractual conditions does natural customer decay not apply?	9
What strategic governance function does an isolated baseline decay rate serve?	9
What operational steps must commercial leadership take before calculating decay?	10
Where are the empirical boundaries of customer asset decay estimates?	10
References	12

MANAGEMENT SUMMARY

A commercial budget records gross investment and never net. Net is what was added after replacing what eroded off the asset base, and that erosion rate is assumed rather than estimated for the asset most commercial teams own: Ewens, Peters and Wang set organisational capital's depreciation at 20% a year and state in a table note that it is not estimated, because the parameter cannot be separately identified from the data. Published rates for brand run from 2.5% to 55% a year and measure different things. This essay works out what follows for a budget, shows by simulation that the obvious way to measure your own rate returns the reason you stopped rather than the behaviour of the asset, and sets out the design that does not.

Keywords: Intangible capital · Depreciation of intangible assets · Organisational capital · Marketing budget allocation · Compounding growth

The go-to-market budget came in at the same number as last year, and nobody argued. A flat line reads as discipline. It was approved in four minutes.

It should not have been. A flat budget on a mature asset does not hold a position; it buys the position back at exactly the rate the position is being lost. That is not a judgement about the team. It is what the arithmetic of any accumulating asset does at a constant rate of spend, and it means a year can add nothing to the asset base while looking, on the page, identical to a year that added a great deal.

Every budget I have seen records gross investment. None has carried the other half of the pair: how much of what the company already owns will wear out over the same twelve months. Net investment is gross minus that erosion, and net is the number that says whether anything accumulated. Its second term has a name in the literature that studies it, a symbol, and, for the asset a commercial team actually runs on, no measured value at all.

Why is maintaining a flat growth budget an active capital commitment?

The arithmetic is the ordinary arithmetic of capital and it has one moving part. This year's stock is what survived from last year plus what was added; what survived is last year's stock less the fraction that decayed; that fraction is the depreciation rate. Hold the spending flat for long enough and the stock stops moving. It settles at the level where each year's addition exactly offsets each year's erosion. That is not a failure, it is the definition of a steady state.

The consequence is sharper than it sounds, because it needs no estimate of anything. Once the stock has settled, the whole investment portion of the spend is replacement. Not most of it. All of it. Whatever share of the budget was building an asset rather than paying for this quarter's output, that share is now holding the asset where it is, and the honest entry against "added to the asset base" is zero.

Revenue can be growing while this is true. Headcount can be growing. What is not growing is what Ingemar Dierickx and Karel Cool called a stock rather than a flow: the part of the commercial machine that makes the next cycle cheaper than the last one. A flat renewal is a decision to stop building it, and it is the only decision in the budget that gets made without anyone saying it out loud.

So how fast does the erosion run? That is where the arithmetic stops and the trouble starts.

What do foundational studies actually report regarding customer decay rates?

The standard reference for what intangible assets are worth in an economy is Carol Corrado, Charles Hulten and Daniel Sichel's 2009 paper in the Review of Income and Wealth. Most later work builds on it, and it says this, in the authors' own words, immediately before the table everyone cites:

Relatively little is known about depreciation rates for intangibles. Based on the limited information available, we made the following assumptions about depreciation rates.

Their assumptions were 33% a year for computerised information, 20% for research and development, 60% for brand equity, and 40% for firm-specific resources: training, management time, reorganisation. On the following page they say where the last one came from: “For firm-specific resources, we averaged the rates for brand equity and R&D.” The depreciation rate for the organisational asset base, in the framework much of this literature rests on, is the arithmetic mean of two other numbers.

That is not a gotcha and it should not be read as one. The paper says what it did, on the page, in a sentence beginning “relatively little is known”. Charles Hulten and Xiaohui Hao are equally direct a year later, capitalising the accounts of a composite pharmaceutical company built from six real ones: faced with published R&D rates running from 10% to 25%, they write that they “somewhat arbitrarily select an amortization procedure near the midpoint of the range”. The candour is at the source. What happens downstream is that the label falls off, and an assumption enters a spreadsheet as a measurement.

Table 1 Every published depreciation rate, and what it was measured on

Rate a year	Half-life	What was measured, and how the number was arrived at
2.5%	27 years	The decay of one year’s consumption experiences inside a consumer’s brand preference. Estimated, from 38,000 households moving between US states, across 238 categories (Bronnenberg, Dubé & Gentzkow)
20%	3.1 years	R&D capital. Assumed: “in the middle of the range of the rates reported in the existing literature” (Corrado, Hulten & Sichel)
20%	3.1 years	Organisational capital: process, coverage, the commercial machine. Assumed: “the [organisational-capital depreciation rate] is assumed to be 0.2 (i.e., not estimated)” (Ewens, Peters & Wang)
33%	1.7 years	R&D capital. Estimated, from the prices paid for 2,004 firms as they left the market (Ewens, Peters & Wang)
40%	1.4 years	Firm-specific resources: training, management time, reorganisation. Assumed: “for firm-specific resources, we averaged the rates for brand equity and R&D” (Corrado, Hulten & Sichel)
55–60%	10 months	Advertising capital. Assumed: “our own interpretation of this literature” (Corrado, Hulten & Sichel; the 55% is the 2016 successor, adopted by Bronnenberg, Dubé & Syverson)

Corrado, Hulten & Sichel (2009), *Review of Income and Wealth* 55(3), pp. 673–674; Bronnenberg, Dubé & Gentzkow (2012), *American Economic Review* 102(6), p. 2474; Bronnenberg, Dubé & Syverson (2022), *Journal of Economic Perspectives* 36(3), reporting Corrado et al. (2016); Ewens, Peters & Wang (2019, revised October 2023), NBER Working Paper No. 25960, Table 1 and its note. Half-lives are the author’s calculation: $\ln 2$ divided by $-\ln(1 - d)$, where d is the annual rate.

Why do customer churn and revenue contraction measure distinct operational assets?


The widest gap in that table is worth pausing on, because the obvious reading of it is wrong.

At one end, 55% to 60% a year for advertising capital. At the other, 2.5%: from Bart Bronnenberg, Jean-Pierre Dubé and Matthew Gentzkow, who tracked more than 38,000 American households across state lines and used the fact that people take their brand preferences with them when they move. Their estimate is not of advertising's persistence. It is, precisely, that "the effect of a given year's consumption experiences depreciates at a rate of ... 0.025 per year", across 238 packaged-goods categories.

A factor of more than twenty, and it is not a disagreement. What advertising buys and what having been used buys are two different sources of the same commercial position, and they hold their value on completely different timescales. Bronnenberg, Dubé and Chad Syverson make the same point about the wider literature ten years later: the debate over the decay rate "has by no means been resolved", and the spread is unsurprising because "different forms of advertising may exhibit different degrees of longevity".

Which puts the practical question somewhere more useful. A depreciation rate is not a parameter inherited from your category. It is a consequence of the mix of things that built the position, and that mix is a decision somebody made. A position assembled out of paid attention behaves like a subscription: it arrives quickly and it leaves quickly. A position assembled out of having been used behaves like a moat, on both counts.

Figure 1 The same asset, left alone, at three published rates



Curves computed as $(1 - d)^t$ raised to the power t , from the rates in Table 1: 2.5% a year (Bronnenberg, Dubé & Gentzkow 2012), 20% (the assumed organisational-capital rate, Ewens, Peters & Wang 2019, revised October 2023), 55% (Corrado et al. 2016 via Bronnenberg, Dubé & Syverson 2022). A disclosed model, not measured data.

Why does commercial management lack external benchmarks for relational equity decay?

Everything above is about brand, because brand is where the measurement effort went. The asset most commercial organisations actually run on is the other one: the process, the pricing architecture, the qualification standard, the account coverage, the working relationships. This literature calls it organisational capital.

The most careful attempt to measure it is Michael Ewens, Ryan Peters and Sean Wang's, held here as NBER Working Paper No. 25960, revised October 2023. They do something clever. Rather than assume what intangible assets are worth, they read it off the prices actually paid when firms left the market: acquisitions, liquidations and bankruptcies, 2,004 firms, and back out the parameters that reconcile those prices with the firms' past spending. It is the best instrument anyone has built for this. The note to Table 1 of the working paper reads:

The [organisational-capital depreciation rate] is assumed to be 0.2 (i.e., not estimated).

They are explicit about why. Two parameters have to come apart and will not: the share of overhead that is investment rather than expense, and how fast the resulting asset decays. In their words, "we cannot separately identify the parameters ... in each capital accumulation process." They then test how much it matters. Raising the assumed depreciation rate from 10% to 30% moves the estimated investment share from 0.18 to 0.40 and leaves the model's fit, in their phrase, "nearly static". A footnote records that organisational capital parameters "have only been estimated in the pharmaceutical industry", which is to say, in Hulten and Hao's six-company composite, in 2006.

So the position is this. Where research and development is concerned there are real estimates, and Ewens, Peters and Wang produce a good one: 33% a year across their sample, against the 23% the national accounts imply for the roughly half of firms they cover at all. Their industry figures I will leave alone, because the standard errors on them run to eighteen points and the spread between industries is not something this data can resolve.

For the asset base a commercial team spends most of its money on, this literature has no estimate and not even a single convention. The finance side uses 20% a year, which its users describe as standard and none of them derived. The national accounts use 40%, which is the average of two other rates. Nobody has brought data to either that can say whether the right answer is 10% or 30%: the difference between an asset that halves in seven years and one that halves in two.

One boundary on that, and it is worth stating plainly, because this essay is about people quoting a rate without checking whose asset it was measured on. What has never been estimated is the parameter these models use: the decay of a stock accumulated out of overhead spending. How fast an organisation loses knowledge it acquired by doing the work is a different question, in a different literature, and it has been studied in production and service settings for decades. Those results are not interchangeable with this parameter: different construct, different data, different asset, and helping myself to them here would be the exact substitution the rest of this piece is about. Where that work lands, and whether any of it travels to a commercial organisation, is a piece I have not written yet.

Why can only internal operating ledgers isolate true baseline customer decay?

The reason sits in Ewens, Peters and Wang's footnote, and it is worth following, because it points straight at what to do.

When the maintaining spend has been roughly constant for a while, the stock it supports settles at that spend multiplied by the investment share and divided by the depreciation rate. The two parameters only ever appear as a ratio. A large share decaying quickly and a small share decaying slowly produce the same stock, the same accounts and the same price at exit. From outside, the level is visible and neither term is.

The ratio comes apart in one situation: when the flow changes sharply, because only then does the stock's own path become visible instead of its resting place. Outside observers know this and hunt for those moments: the "going dark" studies in marketing are exactly that, and they are why anyone believes brand assets decay at all. But an outsider can only find an interruption somebody else chose. A company can design one.


That is the whole asymmetry, and it is not small. It is the difference between an event study and an experiment.

Why does conventional cohort reporting conceal underlying decay mechanisms?

Here is the version of that test most companies have already run by accident. A programme is cut. The number it was supporting drifts down over the following year. The decline, measured against where things stood before the cut, is taken as the asset decaying.

Before recommending anything, I ran that procedure against data with no decay in it whatsoever: a quarterly series wandering around a constant average with the stickiness commercial series have, and an asset that never depreciates by a single percent. Twenty thousand runs, twelve quarters before the stop and eight after. The script is filed with this piece's claim ledger, as the other instrument tests on this site are.

Figure 2 What the stop test reports when there is nothing to report



Simulation, 20,000 runs per design. An AR(1) quarterly series, persistence 0.3, standard deviation 15% of level, with a true depreciation rate of exactly zero; twelve quarters before the stop and eight after. Script: instrument-test-decay.py, filed with the claim ledger.

Cut the programme after a strong run: the ordinary case, because a thing that is going well is the safest thing to stop funding, and the procedure reports a median 6.8% a year of depreciation in an asset that is not depreciating at all. That is larger than the entire published estimate for consumer brand preference. Cut it after a weak run instead and the same procedure reports minus 6.6%: an asset apparently growing while nothing whatsoever is happening.

The naive test does not measure the asset. It reads back the reason you stopped. Both numbers are regression to the mean in a capital-theory costume, and neither has anything to do with depreciation.

Scheduling the stop in advance fixes the bias: the median lands at 0.8%, but leaves an interval running from minus 11% to plus 11%, wide enough to contain every published rate in the table above. And a comparison group is not automatically the answer either. An arm that shares none of your market's shocks removes the bias and makes the interval worse, from minus 17% to plus 15%.

What core econometric baseline survives rigorous retention modeling?

Three conditions, and all three are needed.

Choose the stop before you look. A date fixed in advance, or a region drawn by lot. The moment the choice depends on how the asset is doing, so does the answer.

Run a matched arm that shares your weather. A comparable region, segment or cohort that does not stop. It has to move with yours for the same outside reasons, or it adds noise rather than removing it.

Declare the window in advance and read all of it. Eight quarters is the minimum and twelve is much better, because precision improves faster with length than with anything else you can buy.

Under those conditions the estimate is unbiased, and what remains is a question of how much of it you need. Against a true rate of 20% a year, with a series that wobbles by about 15% of its level from quarter to quarter, two paired stops over eight quarters reach 96% power, and one pair over twelve quarters reaches 98%. At 25% wobble it takes four pairs over eight quarters, or two over twelve. At 35% it takes four pairs over twelve quarters, and below that the honest output is a bound rather than an estimate.

Two paired stops over two years is a real but ordinary cost: one region, one segment, planned like any other measurement. Four pairs over three years is a programme. And these are best cases: the simulation treats the metric you watch as though it were the asset itself, where in practice it is a noisy, lagging function of the asset, which costs precision again.

Under what specific contractual conditions does natural customer decay not apply?

The serious objection is that maintenance may not be separable from operation. A machine left idle rusts. A qualification standard is maintained by being applied, at no cost beyond doing the work that applies it. If commercial assets are largely maintained as a by-product of use, then there is no missing budget line, the erosion term is already inside operating cost, and this argument has imported a factory metaphor and built on it.

That is true of some assets and false of others, and the split is observable. Does ordinary work exercise the asset? A pricing architecture used in every deal is maintained by use. A partner ecosystem, a documentation set, a data model, a category position, the market's understanding of what you are for: these are exercised by a fraction of the work at best and by nobody at worst. The objection does not refute the arithmetic. It locates where the erosion term is near zero, which is worth knowing before spending two years measuring it.

It also stops holding where the asset is embedded in the product, so that no interruption can be isolated; where the series is noisy enough that the design above cannot resolve anything useful; and wherever stopping would be irreversible or would cost a customer something. This is a measurement, not a licence.

There is a milder version of the same problem, and it is the reason this whole quantity went missing in the first place: what gets measured is what some system already has a field for. It is the same mechanism that leaves the most expensive interval in a funnel unrecorded. This confusion is identical to asking what is CAC?, where blended acquisition metrics pool maintenance marketing with incremental demand. No accounting system has a field for the erosion of something it was never allowed to capitalise.

What strategic governance function does an isolated baseline decay rate serve?

Notice what the depreciation rate is not for. It is not for pricing a maintenance line, because that calculation needs the size of the stock, and the size of the stock is exactly what cannot be measured without already knowing the rate. The steady-state result earlier needed no parameter at all, which is why it was the one worth stating.

What the rate governs is time: how long a change in the maintaining spend takes to arrive in the asset. And it governs it symmetrically, in both directions, which is the genuinely useful part.

Table 2 How much of a change in spending has reached the asset

Depreciation rate	After 1 year	After 3 years	After 5 years	After 10 years
2.5% a year	3%	7%	12%	22%
20% a year	20%	49%	67%	89%
55% a year	55%	91%	98%	100%

Author's calculation from the capital accumulation identity: the share of a permanent change in the maintaining flow that has arrived in the stock after t years is $1 - (1 - d)^t$ raised to the power t , where d is the annual rate. Rates as in Table 1.

Read along a row, that is why a slow-decaying asset is impossible to defend in a quarterly review: at 2.5% a year, a permanent increase in spending has delivered 3% of its eventual effect by the time the first annual review arrives. Read as reassurance, it is also why such an asset does not vanish because you looked away for a year. Read down a column, a fast-decaying asset does the opposite: most of the benefit inside twelve months, and most of it gone twelve months after you stop.

Durability and speed are the same parameter with opposite signs. A team choosing what to build is choosing which of the two it wants, and it is usually choosing without noticing.

What operational steps must commercial leadership take before calculating decay?

Two things, and one of them is free.

Put the second line in the budget. Against every item, two entries instead of one: what this replaces, and what this adds. For a mature asset carried at a flat number, the second entry is zero, and writing the zero down is the entire point. It turns an invisible decision into an argued one.

And you already know something about your own rate, from the other direction. The table above runs backwards. If a position took five years of sustained effort before it visibly paid, the adjustment was slow, and a slow adjustment is a low depreciation rate, which means the asset is durable, that it should never have been judged on quarters, and that if it is ever broken it will take five years to get back. If it arrived within a year of the spending starting, it is at the fast end, and it will leave on the same schedule. The build time you already lived through is evidence about the decay rate, and for most companies it is the only evidence they will ever have.

Neither of those is a number you can put in a model, and that is the honest state of the question. There is no importable rate. There never has been. The one place a rate can be established is inside a company willing to stop something on purpose and watch carefully for two years, and very few will. The ones that do will know something about their own business that the entire published literature does not. That asymmetry: what an organisation can see about itself that an outside observer cannot: is also what my doctoral research is built on.

Where are the empirical boundaries of customer asset decay estimates?

Boundary. The stock-versus-flow model is a decision lens, not a universal depreciation schedule. Classify a spend stream only after testing whether its output survives the next cycle.

Evidence base. The analytical frame also draws on these additional sources: Dierickx and Cool 1989; Hulten and Hao 2008. The links identify the exact works; they support the mechanisms and boundary conditions discussed here, not every claim in isolation.

- Bronnenberg, B. J., Dubé, J.-P., & Gentzkow, M. (2012). The evolution of brand preferences: Evidence from consumer migration. *American Economic Review*, 102(6), 2472–2508. <https://doi.org/10.1257/aer.102.6.2472>
- Bronnenberg, B. J., Dubé, J.-P., & Syverson, C. (2022). Marketing investment and intangible brand capital. *Journal of Economic Perspectives*, 36(3), 53–74. <https://doi.org/10.1257/jep.36.3.53>
- Corrado, C. A., Hulten, C. R., & Sichel, D. E. (2009). Intangible capital and U.S. economic growth. *Review of Income and Wealth*, 55(3), 661–685. <https://doi.org/10.1111/j.1475-4991.2009.00343.x>
- Dierickx, I., & Cool, K. (1989). Asset stock accumulation and sustainability of competitive advantage. *Management Science*, 35(12), 1504–1511. <https://doi.org/10.1287/mnsc.35.12.1504>
- Ewens, M., Peters, R. H., & Wang, S. (2019). Measuring intangible capital with market prices (Working Paper No. 25960, revised October 2023). National Bureau of Economic Research. <https://doi.org/10.3386/w25960>
- Hulten, C. R., & Hao, X. (2008). What is a company really worth? Intangible capital and the “market to book value” puzzle (Working Paper No. 14548). National Bureau of Economic Research. <https://doi.org/10.3386/w14548>

END OF ESSAY

HOW TO CITE THIS ESSAY

Isoglu, S. (2026, August 1). Every growth budget is a gross number. Sinan Isoglu.
<https://isoglu.com/insights/journal/every-growth-budget-is-a-gross-number/>

KEEP READING

Growth that compounds

The European failure number nobody quotes

<https://isoglu.com/insights/journal/the-european-number-nobody-quotes/>

Growth that compounds

The thirty-percent rule for sales and marketing

<https://isoglu.com/insights/journal/the-thirty-percent-rule-for-sales-and-marketing/>

Growth that compounds

Why growth compounds, and activity doesn't.

<https://isoglu.com/insights/journal/why-growth-compounds/>



Sinan Isoglu, MBA (Quantic)

Commercial growth leader, lecturer and doctoral researcher