

Reading the *numbers*

How revenue capacity is estimated from published accounts, how breakeven moves with gross margin, and where the method stops working.

What this note is for

Part three of this series rests on a number almost no business publishes and most have never calculated: revenue capacity, the turnover a company's current people, plant and organisation could support if everything sellable were sold. Set actual revenue against it and you have utilisation, which is the horizontal axis of the strategy map and, more usefully, the difference between a business whose next growth is already paid for and one whose next growth needs a cheque.

This note sets out how that estimate is built, why gross margin has to be normalised out before any two businesses can be compared, and where the method stops working. It is written so you can run it yourself, on your own accounts or a competitor's. It is a practitioner's calibration rather than a law of nature: these are where the needle usually settles, not where it must.

STEP ONE

Four readings, triangulated

No single calculation gives you capacity. Four do, read against each other, and the value of the exercise is as much in the disagreements between them as in the answer.

A notional breakeven from the P&L. The turnover at which the fixed cost base is covered. In most businesses this sits near 40 per cent of estimated capacity, which is the anchor the rest of the method hangs from.

A well-run benchmark of around 70 per cent. Where competent operations in most sectors settle before their next wave of investment. A business materially above it is close to its ceiling; one materially below it is carrying capacity it is not selling.

A gross profit per head band. For mid-market businesses in developed markets this commonly falls between \$125k and \$150k. Its job is to check the inputs, not just the answer. A reading far outside the band means one of two things, and it is worth establishing which before going further. Either the business holds an asset advantage worth naming, brand, distribution reach, a technology position, or one of the inputs is wrong. Three input errors account for most cases. The first is how people were counted: full-time equivalents against headcount, and whether contractors, agency staff and outsourced labour are in or out. The second is how gross profit was struck, since businesses differ on what sits in cost of sales, and a services company that puts delivery labour above the line will not compare with one that puts it below. The third is the fixed and variable split behind the breakeven calculation, where treating a semi-variable cost as fixed moves the breakeven point and everything that hangs from it. Fix the input before reaching for the explanation.

The most restrictive bottleneck. The least mechanical of the four and the most important. Something caps what the business can actually deliver, and it is rarely the largest item on the capital register. The register lists what the business owns; the constraint is usually whatever it cannot quickly add. In practice that tends to be one of a small number of things: a licence, accreditation or regulatory cap; a single qualified team or an individual whose sign-off everything passes through; a shift pattern or opening hours; a supplier with a long lead time or a single source; or a step in the process that everything else queues behind. A factory can run a third shift far more easily than it can find another

twenty accredited engineers. The practical test is to ask what would have to change for the business to take on 20 per cent more work next month, and the answer is your bottleneck. It matters because capacity added anywhere else is capacity you cannot sell, and if the arithmetic and the bottleneck disagree, back the bottleneck.


Worked on the Wrigley decade analysed in part two, the triangulation held: gross profit per head of \$173k read correctly as an asset advantage, brand and distribution reach, rather than as a business running unusually hot.

STEP TWO

Normalising for margin

The four readings compare badly between businesses of different margin structure, and the correction is arithmetic rather than judgement. At breakeven, revenue multiplied by margin equals fixed cost. So breakeven utilisation is fixed cost divided by margin multiplied by revenue capacity, and if you hold the cost and capacity structure still, breakeven utilisation is exactly inversely proportional to gross margin. That relationship is arithmetic and holds whatever numbers you put into it.

The numbers below are a different matter. A 40 per cent breakeven on a 50 per cent gross margin is my working midpoint, drawn from the businesses I have looked at rather than from a large sample, and the margin steps either side are round numbers chosen for legibility. Treat the figures as directional and the shape as reliable: as margin falls, the overhead a business can carry falls with it, and faster than most people expect.



Illustrative, on a 40 per cent breakeven at a 50 per cent margin. The shape is the point: at the low-margin end, the same overhead structure stops being affordable.

That last row is the finding, and it is worth stating carefully. On these assumptions, a business on a 20 per cent margin carrying a 50 per cent margin business's fixed costs would be at or beyond breakeven at full capacity, which is to say it has no viable operating range at all. Change the assumptions and the exact figure moves; the direction does not. Margin does not merely shift the breakeven point, it governs how much fixed cost a business can afford to build in the first place. Which is why distributors and contract manufacturers run lean on marketing, management layers and central function, while software and branded businesses carry all three, and why comparing their utilisation without this correction tells you very little.

Turn it round and it becomes a test. At a 40 per cent breakeven, affordable fixed cost runs at roughly 8 per cent of capacity revenue on a 20 per cent margin, 20 per cent on a 50 per cent margin, and 32 per cent on an 80 per cent margin. A business carrying materially more than its margin affords is not under-utilised. It is over-built.

The general form, and a working shortcut

Splitting revenue capacity into its asset turnover and operating capital components gives the fuller statement. Breakeven utilisation equals the fixed-cost intensity of the capital base, divided by margin multiplied by asset turnover. The denominator is the return that capital generates at full utilisation, so the formula reads plainly: *breakeven capacity is the fixed-cost intensity of your capital, divided by the return that capital earns when fully worked*. The margin-only version above is the special case where asset turnover is held constant.

Half of that is calculable from published accounts. Return at full utilisation comes out at roughly 43 per cent for Zendesk in 2021 and for ServiceNow in 2024, against about 25 per cent for Capita and 20 per cent for Serco. The outsourcers turn their balance sheets over far faster, 1.3 and 1.8 times against roughly 0.55 for the software pair, and it does not save them: margin dominates the result.

The other half, fixed-cost intensity, is not available publicly, and it is not an independent gap. For a services business, fixed cost is dominated by central overhead, which is the same unknown as the split between central and delivery headcount. Populate one honestly and you must populate the other the same way, or the ambiguity has moved rather than gone. This is where an outside-in estimate has to become an inside conversation.

THE SHORTCUT

A margin-neutral reading

For working purposes the formula collapses usefully, because margin is already inside gross profit. Utilisation approximates to the assumed breakeven multiplied by gross profit over fixed costs. On a 40 per cent breakeven, the ladder runs as follows.

GROSS PROFIT COVERS FIXED COSTS	IMPLIED UTILISATION	WHAT IT MEANS
once	~40%	Breakeven. A survival problem, not a strategy problem.
1.4 times	~55%	Efficiency band. Sell the capacity you already own before buying more.
1.75 times	~70%	The investment trigger. The agenda shifts from whether to grow to how.
twice	~80%	The practical ceiling. Sustained operation here raises its own questions.
2.5 times or more	past 100%	Off the scale, which is a finding rather than a failure. See below.

Treat these multiples as a cross-check on the other readings rather than a replacement for them. They are quick, they travel between margin structures, and they will get you close enough to know which conversation to have.

Turning the bands into a decision rule

The calibration becomes useful the moment decisions are attached to it. Take a business turning over £10m at 70 per cent of capacity. Total capacity is roughly £14.3m, and the trigger has been reached. Invest to take capacity to £20m and utilisation re-enters at 50 per cent, still comfortably above the breakeven floor, with room to grow into.

A board can adopt that as policy. Invest at 70, re-enter at 50, never plan below 45. The value is not the precision of the numbers, which will be wrong in the third decimal place and will not matter. It is that the trigger is agreed before the year that provokes it, which turns capacity investment from a reaction to a good year into a decision whose conditions were set in advance.

LIMITS

Where it breaks

The method assumes capacity is broadly proportionate to the capital and people on the balance sheet. Two situations break that assumption, and both are informative rather than fatal.

Capacity outside the perimeter. Where a business has deliberately placed capacity on someone else's balance sheet, as Coca-Cola does with its bottlers and Apple with its manufacturing partners, the reading goes off the scale. Measure the system and it comes back inside the normal range; measure the entity alone and it reads past 100 per cent. The gap between the two measures how much capacity the model has pushed outside the perimeter.

Capacity that never hit a register. Brand strength, distribution position, ecosystem, technology. Expensively built, quietly compounding, and invisible to any calculation that starts from plant and payroll.

One caution attached to both. A capacity ceiling is a sector benchmark, not a law of physics: what a typical business of that type achieves with typical assets. Exceeding it says you hold something the reference group does not, and the question is whether you can name it. Where there is a nameable, durable asset behind the number, the reading is capability. Where there is not, the likelier explanation is a business trading beyond what it can sustain. Same number, near-opposite diagnosis.

A closing word on the calibration. The thresholds here are a working set, built from the businesses I have analysed rather than a broad sample, and I expect them to move as that set grows. The method does not depend on them: the arithmetic holds whatever midpoint you adopt, so if your sector's breakeven sits elsewhere, substitute it and everything downstream still works.

Which is the honest summary of this note. The arithmetic will get you a defensible number and a first view of which agenda leads. What it cannot tell you is whether that number is capability or exposure, and that question is answered inside the room, not from the accounts.

Companion to *Bought or Learned*, part three of The Value Creation Imperative. Wrigley from Wm. Wrigley Jr. Company annual reports (2007; SEC EDGAR). Zendesk and ServiceNow from filed 10-K data. Capita and Serco from FY24 statutory accounts. Bands and thresholds are a practitioner's calibration; figures are approximations.