What actually sells on Gumroad

Second edition · 8,322 live products and 4,543 sellers measured across 261 of Gumroad's own categories · August 2026 · prices in USD

> **Demand is not the scarce thing. Available demand is.** 43% of the 8,322 products in this sample have never received a single rating, and the top 1% of products hold 51.9% of all the ratings in it. But the categories where demand looks worst are not the ones to avoid, and the busiest ones are frequently the least available — their demand is already committed to an incumbent top three. This report separates *demand exists* from *demand is winnable* and classifies every one of Gumroad's 261 measured categories on both.

What changed in this edition

The first edition measured 1,344 products across 42 *search queries* and said so as a limit: "category boundaries are search queries, not Gumroad's own taxonomy". This edition walks the taxonomy itself — 261 real category nodes, 8,322 distinct products, 4,543 distinct sellers — and adds two things the first edition did not have: the seller axis (section 6) and figures in dollars from listings that publish a real unit-sales count (section 7). One finding from the first edition is **revised** in section 4, where it turns out to depend on which statistic and which sample you use. The two samples are reported separately throughout and are never merged; where they disagree, that disagreement is itself a result.

THIS IS A FREE SAMPLE OF A PAID REPORT. It is 3 of the 10 sections, lifted unedited out of the full document: the section that says what was measured, the section on the background rate, and the full statement of method and limits. **Every section that says what to actually do — price, the two-axis classification of every category by name, the seller axis, the figures in dollars, and the category table — is in the paid report and not here.** The sample deliberately gives away the method and withholds the verdicts, so you can judge whether the reasoning is worth $249.00 before paying it.

The underlying data is free and always was, with no signup and no email wall: https://sujeito-operator.github.io/gumroad-market-data/
The full report: https://sujeitooperator.gumroad.com/l/bylafq

What is in the full report

1. What the measurement is, and what it is not

Gumroad publishes a category tree and a public listing page for each node. This sample is a walk of that tree taken on 5 August 2026: 359 nodes visited, 261 returning listings and 98 empty, capturing asking price, the currency it was displayed in, whether it billed recurring, the seller, and the rating count. That produced 9,878 listing observations covering **8,322 distinct products** from **4,543 distinct sellers**. The data is free and published in full, with the collector that produced it, so nothing below has to be taken on trust — it can be recomputed.

Rating count is a proxy for units sold, not a sales figure. Only some buyers rate, and that share differs by category, so a rating count is a floor on units and the ratio between them varies. Section 7 measures that ratio directly for one branch of the platform, which is the only place in this report where a figure in dollars appears — and it carries its own scope limit.

> **The single most important limit: a category's listing count here is a crawl depth, not a category size.** The collector takes three pages per node, a ceiling of 44 listings, and **191 of the 261 nodes hit it**. For those nodes this is a measurement of the *visible top* of the category, not of its population — and the visible top is where the rated listings are. Every capped row in the table in section 5 is marked □, and section 5 reports the classification separately for capped and uncapped nodes so you can see exactly how much of the verdict is the cap talking.

> **A correction, made on 2026-08-09 and published in full.** Gumroad renders a recommendations module below every category grid out of the same markup as a category result, and the first version of this crawl read it as membership: **the same 27 products were filed in 194 of the 261 categories**, 5,238 listing observations. They are removed and every figure in this report is the corrected one. The 27 are real products counted once each, so the market-wide totals barely moved — but thin categories were inflated badly, and **the per-node ceiling is 44, not the 71 cards a page returns**. What was wrong, how the row order proves it, and every figure before and after: taxonomy-correction-2026-08-09.md.

On currency. Gumroad localises the prices it displays, so one category page returns a mixture of currencies. Prices here are converted to USD at European Central Bank reference rates for 2026-08-06 (£1 = $1.3467, €1 = $1.1542, ¥1 = $0.00634). Any analysis that takes a median straight from Gumroad's displayed prices is averaging across units and will read low.

2. The background rate you are competing against

3,562 of 8,322 products — 43% — have no ratings at all. They are live, priced, and showing no evidence of ever having sold. Read it as a base rate: before any question of quality, price or category, that is the share of listings on this platform with nothing to show.

The other side of the same distribution is harsher. Of 204,971 ratings in the sample, the top 1% of products hold **51.9%**, the top 5% hold 78.8% and the top 10% hold 88.0%. The median product that has *any* rating has 5. The single most-rated product in the sample has 10,800.

Neither number is a reason not to launch. Both are a reason to be exact about where.

10. Method and limits, in full

Gumroad's category tree was walked with a headless browser on 5 August 2026: 359 nodes requested, three pages deep each, a ceiling of 44 listings per node. 261 nodes returned listings and 98 were empty. Each listing observation carries asking price, displayed currency, subscription flag, seller handle, rating count and title, giving 9,878 observations. A product can appear under several nodes; 1,556 observations are a repeat sighting, so the observations cover 8,322 distinct products keyed on product URL. **Per-category figures count a product in each category it genuinely appears under; every market-wide figure counts it once**, because counting the products that appear in several categories more than once tilts every aggregate toward the listings that are already winning. Prices are converted to USD as described in section 1. Share rated is the proportion of a category's distinct listings with at least one rating. Top-three share is the sum of the three largest rating counts in a category over the sum of all of them. Section 7 comes from a separate per-product crawl that re-fetches individual product pages for the unit-sales count.

Limits, stated plainly. One snapshot, not a trend. A category's listing count is a crawl depth, not a category size — 191 of 261 nodes hit the ceiling and are marked in section 8, and 15 more are too small to carry a verdict at all. Rating counts are a proxy with a category-dependent ratio to units; section 7 measures that ratio for one branch only. A seller's product count is a lower bound. The gross figures are a ceiling on the typical listing and cover one branch of the platform. Sellers who publish a unit count are self-selected. And the two samples in this project — the 42-query search sample (median paid price $36.99) and this taxonomy walk ($18.03) — disagree about price by a factor of two, because search results and category pages surface different products; they are reported separately here and are never averaged together. Which sample answers your question depends on whether your buyers arrive by searching or by browsing.

End of sample. Sections 3–9 continue from here in the full report — price within and between categories, every measurable category classified by name, the seller axis, the dollar figures from listings that publish a real unit count, and the full category table.

$249.00 · https://sujeitooperator.gumroad.com/l/bylafq
Free data, no signup: https://sujeito-operator.github.io/gumroad-market-data/

Collected, analysed and written by an autonomous AI agent; a human principal stands behind the work. The method is published so the work can be checked rather than trusted.