



KEN RESEARCH · METHODOLOGY

B2B software didn't hide its prices. It added Contact sales

The data, the methods, every number and every limitation behind the study.

Cristian Frunze

Founder of Ken

ABOUT THE DATA

This study uses Ken internal data: the company, people and web records Ken uses for sales and outreach, reused by Cristian Frunze for this study.

Aggregates only. Any group smaller than 50 people is suppressed, and no individual is named.

Figures 1 to 3 are the study page's charts, drawn from the same numbers. The methodology after them keeps the full tables.

SUMMARY

B2B software didn't hide its prices. It added Contact sales

Prices stayed on the page: 72.1% of B2B software pricing pages showed one in 2024-2026, against 75.8% in 2014-2016. What spread is Contact sales, offered on 60.3%, up from 44.4%.

70.8% to 79.1%

of classified pricing pages showed at least one price in every year from 2014 to 2026. We had expected a rise of at least 10 points.

20.2% to 32.3%

of pricing pages showed plans next to a Contact sales tier, 2014-2016 to 2024-2026. Pages with prices and no Contact sales fell from 55.6% to 39.7%.

11.4%

of consecutive observed year pairs switched between showing prices and Contact sales only, about evenly in each direction (52.9% toward hiding).

+3.0

log points of extra headcount growth over eight quarters after a company switched to Contact sales only, against matched peers (95% interval -8.8 to +17.1).

Data and method


Data	Pricing pages from Ken's web archive, one crawl a year, for B2B software companies in Ken internal data (the company records and work histories Ken uses for sales and outreach, reused by Cristian Frunze for this study). Funding from SEC Form D filings.
Sample	10,000 B2B software companies and 10,610 pricing pages read. 252 classified company-years in 2014-2016 and 1,800 in 2024-2026.
Date range	Yearly crawls from June 2014 to September 2026. Headcount from 2015 Q1 to 2026 Q1.
Method	Pre-registered. A language model coded each page, with a second pass added after checks against AI annotators on 160 and 80 pages. Yearly shares with 95% intervals, year-to-year switches and a matched event study. Headline numbers recomputed by a second, independent analysis.
Excluded	• Company-years with no pricing page found, or too few captured pages to tell • Pricing pages whose prices are drawn by JavaScript (status unknown) • Groups and outcomes with fewer than 50 companies
Privacy	Aggregates only. Any group smaller than 50 people is suppressed, and no individual is named.

FIGURES

The study's charts

FIGURE 1

Prices stayed on three in four pricing pages, and Contact sales spread




Share of classified pricing pages each year that show at least one price, and that offer "Contact sales" (as a tier or on its own). A page with plans and a Contact sales tier counts in both. Bands are 95% intervals; 53 to 787 companies a year.

Source: Ken internal data (Ken's archive of company websites and the company and career records Ken uses for sales and outreach), analysed by Cristian Frunze.

FIGURE 2

No version of the data shows the 10-point rise we expected




Change in the share of pricing pages that show a price, 2014-2016 to 2024-2026, in percentage points, with 95% intervals. The first rows are cuts of the data; the "Corrected" rows adjust for coding errors using the annotators' labels.

Source: Ken internal data (Ken's archive of company websites and the company and career records Ken uses for sales and outreach), analysed by Cristian Frunze.

FIGURE 3

Each period, more pricing pages pair plans with a Contact sales tier



Mix of classified pricing pages by period. Each bar adds up to 100%. Classified company-years per period: 252 (2014-2016), 1,339, 1,852 and 2,587 (2023-2026).

Source: Ken internal data (Ken's archive of company websites and the company and career records Ken uses for sales and outreach), analysed by Cristian Frunze.

The data, the methods and every number

This document, from October 2026, explains how the numbers in the article were made: the data, the rules, the tests and their limits.

1. Question and pre-registration

The question: what share of B2B SaaS companies publish their prices instead of saying "Contact sales", from 2014 to 2026, and do companies that hide or reveal their prices grow differently afterwards?

Three hypotheses, the method and what would prove each one wrong were pre-registered on 2026-10-16, before any data was pulled:

1. Price transparency rose: the share of pricing pages that publish at least one price is at least 10 points higher in 2024-2026 than in 2014-2016, and most of the rise is pages that publish plans next to a "Contact sales" tier.
2. Switching is rare (under 10% of consecutive observed years), runs both ways, and among companies with more than 200 people, switches to hiding outnumber switches to showing.
3. Growth leads hiding, hiding does not lead growth: hiders grew faster before the switch than companies that kept their prices, and after matching, 8-quarter headcount growth after the switch differs by less than 5 points.

Every later change is logged with its reason in the study's decision log and listed in section 10. The biggest one: the pre-registered page coder turned out to miss most "Contact sales only" pages, so a second coding pass was added (section 3.3). Results are reported for both codings.

2. Data

2.1 Companies

The sample is the Ken Research homepage panel (built for the Ken Research study

Half of B2B software homepages now say AI. Did it pay off?

): 10,000 B2B software companies drawn at random (fixed seed) from company records in Ken internal data, Ken's dataset of public professional profiles, work histories and company records. Industries: software development, technology, information and internet, security, data and other software. Headquarters in the US, the UK, Canada, Australia, Ireland or New Zealand. Founded 1950-2018, 11-5,000 employees on the current record. A company enters the study in 2014 or its founding year, whichever is later: 122,115 company-years.

This is a sample of companies that exist today. Companies that closed before 2026 are not in it, so the early years describe today's survivors, not all software companies of 2014.

2.2 Pricing pages

Ken's web archive, one crawl per year (the homepage panel's mid-year crawl): June 2014, July 2015, June 2016, June 2017, June 2018, June 2019, June 2020, June 2021, July 2022, June 2023, June 2024, June 2025 and September 2026.

For each crawl we listed every captured URL on each company's site: its registrable domain, or only the host itself for the 149 panel companies whose site is a subdomain of a bigger site. The lists came from the archive's URL index.

From each site's URL list we took pricing-like paths:

- **Standard:** /pricing, /plans, /plans-and-pricing, /pricing-plans, /prices, /price, /subscription(s) and close variants, with an optional language prefix and file ending (the rule of the Ken Research study *What does a normal B2B SaaS pricing page look like?*).

- **Unusual:** the same words deeper in the path or in a subdomain (`/product/pricing` , `/pricing/teams` , `pricing.example.com`), plus `/buy` , `/purchase` , `/editions` , `/packages` , `/compare-plans` , `/licensing` , `/get-a-quote` and `/request-pricing` . Blog, news, docs, help, research and similar sections are excluded.

The two best candidates per company and crawl (standard before unusual, then pricing before plans before other words, then shorter paths) were fetched from the crawl archive with ranged reads: 10,610 pages. Every page's HTML was turned into text the same way for every year.

The crawl also gives, per company and year, the number of distinct HTML pages captured on the site (crawl depth). We use it to tell "no pricing page" from "not seen". Coverage grows a lot over time: in 2014-2016 the crawls captured 10 or more pages for 562 to 914 of the panel companies a year, against 3,854 to 4,345 in 2024-2026, and the pricing status is known for 84 companies a year on average, against 600. The checks in section 6 showed this is the crawls, not our reading of them.

Panel companies per yearly crawl (log scale): with any captured page, with 10 or more captured pages, with a pricing page found and with a known pricing status.

2.3 Growth and funding

- **Headcount:** the company headcount history (built for the Ken Research study *Did AI kill the SDR role?*): quarterly headcount per company from 2015 Q1 to 2026 Q1, rebuilt from the start and end dates of people's roles in Ken internal data. 7,793 of the 10,000 companies have people in it. The rebuilt numbers are smaller than real headcount (median about a third of the employee count on the company record), because Ken internal data covers six countries and not every role is linked to a company. They are good for comparing companies with each other and over time, not as absolute sizes. Only about 1% of the companies have more than 200 people in it.
- **Funding:** SEC Form D filings matched to the company records (built for the Ken Research study *When should you hire your first salesperson?*): 1,277 offerings for 449 companies. Form D covers US private raises only, so it is a sparse signal.

3. Coding the pages

3.1 First pass (pre-registered)

Each fetched page was coded by a language model with the fixed schema and prompt of the pricing page study, unchanged: a small language model, a forced structured output and the first 14,000 characters of the page text. The schema records whether the page is a pricing page and every plan with its printed prices, free plan and "Contact sales" flags. The second candidate was coded only when the first was not a pricing page, an HTTP error or an empty page. 8,975 pages were coded.

3.2 Validation of the first pass

We drew 160 coded pages at random from the even-year crawls 2014-2024 (26 or 27 per year). Four separate AI annotators read them, 40 each, without seeing the coding and without model calls, and filled in the same fields (is it a pricing page, does it print a price, a free tier, a "Contact sales" tier, blank price slots). Their brief counted a page as a pricing page when it presents prices, plans or packages, even if the prices say "contact us".

The first pass was precise but missed many pricing pages. Same status as the labels on 74.4% of pages. It called 15 of the 19 "contact sales only" pages "not a pricing page": a `/pricing` page that only says "contact us" has no plans to record. It also called some priced pages "contact sales only" (precision 0.36 for that status). The miss was larger on old pages, so it could change the trend: on the labeled pages of 2014-2018 the coder found prices on 89.3% of the pages it classified (25 of 28), against 70.7% by the labels (29 of 41); for 2020-2024 the figures were 75.0% (27 of 36) and 84.4% (38 of 45).

3.3 Second pass (added)

A second pass re-read every page the first pass had coded as not a pricing page, contact sales only or JavaScript-hidden (5,130 pages, including the uncoded fallback page of company-years whose first page lost its status). It uses the annotators' field definitions plus one sentence: a page whose purpose is to say what the product costs or how to get a price is a pricing page, even when it lists no plans or prices. Pages the first pass coded as showing prices kept that coding, because its "publishes a price" precision was 1.00. Same model, same runner, same text.

With the second pass, status agreement on the first sample rose to 85.0% (from 74.4%), with "contact sales only" precision 0.95 and recall 1.00. That sample shaped the prompt, so we tested it again on 80 fresh pages from the second-pass set (40 from 2014-2018, 40 from 2020-2026), labeled by four new AI annotators whose brief carried the second pass's broader definition. Same status on 66 of the 80 pages (82.5%; the first pass on the same pages: 46 of 80), "pricing page" precision 0.95 and recall 0.79, "contact sales only" precision 0.89 and recall 0.86. The remaining errors differ by era: "contact sales only" precision 0.62 and recall 0.71 in 2014-2018, against 1.00 and 0.91 in 2020-2026. Section 7 takes this into account.

4. Status rules

Per page:

- **Shown:** prints at least one plan price (a free plan counts), and no plan says "Contact sales".
- **Partly shown:** prints at least one plan price, and at least one plan says "Contact sales", "Custom" or similar.
- **Contact sales only:** a pricing page that prints no price.
- **JavaScript-hidden:** plan cards with blank price slots, or an empty JavaScript app shell. The status is unknown and counted on its own.
- **Not a pricing page.**

Per company-year: the first candidate if it is a pricing page, otherwise the second. A pricing-URL capture with no text at all counts as JavaScript-hidden. With no pricing page found, the company-year is **no pricing page** when that year's crawl captured at least 10 distinct HTML pages on the site, and **not observed** otherwise.

Classified company-years are shown, partly shown or contact sales only. "Shows a price" means shown or partly shown. "Offers Contact sales" means partly shown or contact sales only; this measure was added after the pre-registration and is reported next to the pre-registered one. All shares are shares of classified company-years. Version 1 uses the first pass only; version 2, the main version, uses the second pass for the pages it re-read.

5. Analyses

1. **Trend.** Per year and per period (2014-2016, 2017-2019, 2020-2022, 2023-2026), the share of classified company-years in each status. Yearly intervals are Wilson 95% intervals; period intervals come from 1,000 bootstrap draws of companies.
2. **H1 test.** The share that shows a price in 2024-2026 minus 2014-2016, pooling company-years within each period. 95% interval from 2,000 bootstrap draws of companies; a company's early and late years stay together.
3. **Transitions (H2).** Pairs of consecutive classified years of the same company, at most 2 years apart. Hiding: shown or partly shown, then contact sales only. Showing: the reverse. Size at the time is the company's headcount in the quarter of the first capture (2015 Q1 for 2014 captures). Added check: pairs whose two captures have identical text gave the coder identical input, so any switch between them can only be coding noise.
4. **Growth after a switch (H3).** An event study with a matched comparison. Event: a company's first switch in a direction. Controls: pairs with the same starting status that did not switch. Exact matching on the year of the

second capture, size at the first capture (1-10, 11-50, 51-200, 201-500, 501+ people), industry group, founding-year bin, prior-growth tercile (headcount change over the 4 quarters before the first capture) and a Form D raise in the 3 years before. The estimate is the switcher-weighted average of within-stratum differences. Outcomes: headcount change (log of 1 + headcount) from the second capture to 4 and 8 quarters later, a Form D raise in the next 2 years, and the site disappearing from later crawls. Intervals from 2,000 bootstrap draws of companies. Pre-trend check: the same comparison over quarters 8 to 4 before the first capture. An outcome measured on fewer than 50 switching companies is not reported.

6. Checks

- **Agreement with the pricing page study (2026).** That study coded the 2026 pricing pages of the same companies from a different text extraction. On the 402 pages in both, the first-pass status agrees on 88.8% and "shows a price" on 96.3%.
- **Validation** (sections 3.2 and 3.3).
- **URL lists.** For 2020, 2022 and 2024 our URL lists agree with the homepage panel's own lookups: they hold a capture for every company whose homepage the panel fetched.
- **"No pricing page" (2026).** Among 2026 company-years we call "no pricing page", 20.7% have a pricing link on the homepage the panel fetched (22.5% with the first pass only); the crawl missed their pricing page. "No pricing page" is therefore overstated, and the headline uses classified company-years only.
- **JavaScript.** Each year, 0% to 5% of coded pricing pages are JavaScript-hidden (status unknown). Among the 1,769 "contact sales only" pages, 9.6% hold price-like strings inside scripts in the raw HTML; a sensitivity analysis counts those pages as showing prices.
- **Unusual URLs.** 35.4% of classified company-years come from an unusual URL, where fewer pages show a price (64.4%, against 79.2% on standard URLs); the headline is recomputed on standard URLs only.
- **Independent recomputation.** An independent re-analysis, written from the study plan alone without seeing our code, recomputed the company-year statuses and the H1 to H3 numbers from the raw codings. All 13 years of company-year status counts and the H1 and H2 numbers match exactly, apart from bootstrap intervals that differ in the first decimal. The H3 numbers match exactly when ties in prior growth are ranked the same way. Its first run also found one malformed coding value that our code had misread; that was fixed before the final run.

7. Robustness and measurement error

The H1 difference (version 2: -3.7) is recomputed for B2B software pages only (-4.0), standard URLs only (-13.4 points, 95% interval -20.4 to -5.1), companies classified in 3 or more years (-5.9), companies classified in both periods (-4.1, 73 companies), JavaScript-hidden pages counted as showing prices (-3.8) and "contact sales only" pages with prices in scripts counted as showing prices (-0.4). The pre-registered coding gives -5.6. No variant comes near the 10-point rise that H1 predicted; on standard URLs alone the share fell.

The coding still makes errors, and they differ by era. A sensitivity analysis reweights each company-year by the label distribution of its coder class in a validation sample: the chosen page's status (shown, partly shown or contact sales only) or, for company-years whose candidates the coder calls not a pricing page, the first candidate's. Labels that are not a pricing page, or JavaScript-hidden, drop out. Each sample has a weakness. The fresh sample covers only pages the second pass re-read. The first sample shaped the second-pass prompt, and its annotators used the narrower pricing-page definition, under which a page that only says to contact sales, with no plans, is not a pricing page. So four calibrations are reported:

Calibration	Shows a price	Offers "Contact sales"
None (version 2 as coded)	-3.7 (-10.9 to +3.8)	+15.8 (+6.8 to +24.3)
Fresh sample (main check); priced pages taken as correct	-11.0 (-21.3 to -0.2)	+20.0 (+9.2 to +31.1)
Matched definitions: fresh-sample cells for contact-only and not-a-pricing-page, first-sample cells for priced pages	-11.5 (-22.0 to -0.8)	+13.4 (-2.7 to +30.6)
First sample only	+5.2 (-4.8 to +16.0)	+5.4 (-11.9 to +22.3)
Both samples pooled	-4.0 (-11.6 to +4.6)	+9.7 (-6.4 to +25.0)

Changes are 2024-2026 minus 2014-2016, in percentage points, with 95% intervals from 2,000 draws that resample the companies and the validation pages within each cell. The cells are small (7 to 68 pages), so this is a check on direction, not a better estimate. No calibration gives the 10-point rise in published prices that H1 predicted. The share that offers "Contact sales" rises in every calibration, by 5.4 to 20.0 points, but three of the four intervals include zero.

8. Verdicts

- H1: rejected.** Version 2: 75.8% in 2014-2016 (252 company-years, 141 companies) and 72.1% in 2024-2026 (1,800; 1,164), a change of -3.7 points (95% interval -10.9 to +3.8). Pre-registered coding: 83.7% and 78.2%, a change of -5.6 points (95% interval -12.0 to +1.4). The "partly shown" share changed by +12.1 points (95% interval +5.2 to +18.8); pre-registered coding: +13.2 points (95% interval +5.6 to +20.4). But there is no rise in published prices for it to explain.
- H2: rejected.** In version 2, 11.4% of 3,486 pairs (1,243 companies) switch (95% interval 10.4 to 12.5), above the 10% bound, and switches run both ways (52.9% are to hiding). Among companies with more than 200 people at the time (57 companies), 12 switches went to hiding and 10 to showing: the predicted direction, on too few switches to carry weight. With the pre-registered coding, 9.8% of pairs switch (95% interval 8.8 to 11.0), under the bound, and the size part cannot be tested (fewer than 50 companies with more than 200 people). Among pairs whose two captures have identical text, 0 of 269 switch.
- H3: rejected** (part a fails, part b holds). (a) Hiders did not grow measurably faster in the year before the switch: +3.0 log points (95% interval -1.0 to +7.0), 182 hiders against 2,317 pairs that kept publishing. (b) After matching, headcount growth in the 8 quarters after hiding differs by +3.0 log points (95% interval -8.8 to +17.1), with 85 switchers and 290 control companies. The pre-registered rule rejects (b) only for a difference of 5 points or more with an interval that excludes zero; this interval includes zero, so (b) holds, but the estimate is imprecise rather than small. Pre-trend: +5.2 log points (95% interval -2.1 to +13.5). Companies that start showing prices: -10.2 log points (95% interval -28.9 to +13.2), 62 switchers. Pre-registered coding: +10.5 log points (95% interval -6.3 to +24.4) after hiding, 59 switchers. Growth is the change in $\log(1 + \text{headcount})$, times 100.

9. Limitations

- Survivors.** The sample is companies that exist in 2026. The early years describe today's survivors; companies that closed are missing, and if they priced differently, the early shares are off.
- Thin early crawls.** The early crawls captured few pages per site, so early company-years come from better-crawled, probably bigger or better-linked sites. The same-companies variant (-4.1) agrees with the headline, but rests on 73 companies.

- **Coding errors.** The second pass still disagrees with the annotators on 14 of 80 fresh pages, and its errors differ by era. Corrected for them, the change in published prices runs from -11.5 to +5.2 points (section 7). The "no 10-point rise" result holds in every version; the sign of a small change does not.
- **Classified company-years only.** Shares leave out "no pricing page" and "not observed" company-years. That assumes the crawl misses pricing pages regardless of what they show; if it misses "contact sales only" pages more often (fewer links to them, for example), those pages are undercounted.
- **JavaScript.** Prices drawn by scripts are invisible in the page text. JavaScript-hidden pages are counted as unknown; counting them as showing prices barely moves the result (-3.8).
- **Candidate URLs.** Pricing pages are found by their URL. A pricing page at an unexpected address is missed, and unusual URLs show prices less often than standard ones (section 6).
- **Headcount.** Rebuilt from work histories in Ken internal data: about a third of the employee count on the company record, from six countries. Only about 1% of companies have more than 200 people in it, so the size part of H2 rests on few switches.
- **Funding.** Form D covers US private raises only.
- **Yearly snapshots.** One capture a year (pairs at most 2 years apart) misses changes that reverse within a year, and coding errors on changed pages can create switches that did not happen.
- **Correlation, not cause.** Matching removes measured differences (year, size, industry, age, prior growth, funding), not unmeasured ones, such as a new product line or a new sales leader, that could drive both pricing and growth. The 8-quarter intervals are wide.
- **Domains.** Companies are followed under today's domain; a company that changed its domain is not observed before the change.

10. Changes from the pre-registered plan

All are logged with dates and reasons in the study's decision log.

1. **Second coding pass** (section 3.3). The pre-registered coder missed most "contact sales only" pages. The headline uses the coding with the second pass; every result is also reported for the pre-registered coding.
2. **Validation and measurement error.** Two validation samples labeled by AI annotators (160 and 80 pages) and the error-corrected sensitivity analysis with four calibrations (section 7) were added.
3. **2026 pages** were fetched as raw HTML and coded from the same text extraction as every other year, instead of reusing the codings of the pricing page study, which serve as a check (section 6).
4. **Primary crawls only.** One crawl per year; the fallback crawls were skipped for time.
5. **Unusual-URL rule** tightened after reading the 2026 candidate paths, before any page was fetched or coded.
6. **Added checks and variants:** the same-text noise check, the event study with only switchers whose page text changed, the H1 variant with companies classified in both periods, and the 2016 coverage check.
7. **Added measure:** the share that offers "Contact sales" (partly shown or contact sales only), reported next to the pre-registered measure. The H1 verdict uses the pre-registered measure.
8. **Small groups:** event-study outcomes measured on fewer than 50 switching companies are reported as counts only, and the H3 verdict scores parts (a) and (b) separately. An exploratory H2 split by relative size (top quarter of rebuilt headcount in the year) was added, labeled exploratory.
9. **Procedure only:** the crawl order was changed twice and the index reads were run in parallel. No result depends on these.

11. Data handling

Codings, labels and page text stay on Ken's servers. Published outputs are aggregates only, and no reported group has fewer than 50 companies.

Ken Research · B2B software didn't hide its prices. It added Contact sales · Cristian Frunze, Founder of Ken