FROM THE RESEARCH BENCH

A vendor page supplied the source for the 5x retention rule

Asked where the figure comes from, an assistant named Bain and Frederick Reichheld three times out of three. No search finds it in either's work.

Sinan Isoglu, MBA (Quantic)
Commercial growth leader, lecturer and doctoral researcher

PUBLISHED
28 July 2026

READING TIME
8 min

WORDS
1,668

https://isoglu.com/journal/a-vendor-page-invented-the-source/

CONTENTS

MANAGEMENT SUMMARY

The claim that acquiring a customer costs five to twenty-five times more than retaining one is quoted everywhere and has no recoverable study behind it. That much has been known since 2005. What is new is the attribution: across twelve counted runs — four assistants, three each — one of them named Bain & Company and Frederick Reichheld in all three of its own, and the wording came from a staffing company's statistics page it cited. The published trail runs somewhere else entirely and passes through neither. Nineteen of thirty-six answers in the wider study carry a provenance error. The piece does not establish how the attribution reached every system that gave it, and it tests API endpoints rather than consumer chat products.

Keywords: customer acquisition versus retention cost · AI search provenance · citation accuracy in generative search · Frederick Reichheld · commercial statistics

Ask an AI assistant how much more expensive it is to win a customer than to keep one, and tell it to cite the original source. One of them named Bain & Company and Frederick Reichheld three times out of three. Neither has published that figure. The attribution came from a virtual-assistant staffing company's statistics page, which the assistant cited, and which it reproduced almost word for word.

That the number itself is unsourceable has been known for twenty years. Where the source came from is the new question — because somebody had to supply one, and somebody did.

The companion case, and why this one is different

A companion piece covers what these systems do when a real paper exists: they return the preprint, carrying figures peer review deleted, and cite them accurately. This is the other half of the same study. Here there is no paper to miss. The question was chosen precisely because nothing sits at the end of the trail, so whatever the systems name as an origin has to have come from somewhere else.

Twelve counted runs — four assistants, three each, web search on, every prompt closing with the same instruction: give the figure and cite the original source. All twelve went to API endpoints; the consumer chat products behave differently. The protocol, the responses and the coding are published with this piece.

Every one returned a figure. None produced a primary source, because none exists. The interesting part is the four different ways they handled that.

- Claude asserted an origin, three times out of three: "The original source for this widely referenced statistic is Bain & Company research conducted by Frederick Reichheld."
- Gemini named a real document that does not contain the claim, on two of three runs: "a 1990 Harvard Business Review article by Frederick Reichheld, cited the cost as 5 times more."
- ChatGPT named no origin at all on two of three, and filled both citation slots from one growth-consultancy blog.
- Perplexity reported the attribution as an attribution on all three — "The original source most often traced for this rule of thumb is Frederick Reichheld's Bain & Company research" — which is a true sentence, and the only construction in the cell that survives contact with the evidence.

Where Claude's sentence came from

It cited four pages, and the sentence tracks one of them: Stealth Agents, a staffing company for virtual assistants, running a page of retention statistics. That page says:

> **The foundational figure here comes from Bain and Company research by Frederick Reichheld: acquiring a new customer costs 5 to 25 times more than retaining an existing one.**

The assistant's cited sentence keeps the attribution intact and swaps the tail: "The foundational figure comes from Bain and Company research by Frederick Reichheld, which established this 5-25x range." On a second run it also carried the page's floor-and-ceiling gloss, that 5x suits transactional consumer

businesses and 25x complex B2B. The retrieval worked. The page is where the origin was manufactured.

Step	What it says	What is actually there
The assistant	Bain & Company research by Frederick Reichheld established the 5–25x range	No such publication has been found
The page it cited	The same attribution, near-identical wording	A statistics page selling virtual-assistant staffing
Reichheld & Sasser, HBR 1990	Named as the origin by one of the four systems, echoed as an attribution by a second	No cost ratio of any kind appears in it
The published trail	—	TARP, Hart/Heskett/Sasser, Tom Peters. No Bain, no Reichheld

The trail that was already published

In Loyalty Myths (2005), Keiningham, Vavra, Aksoy and Wallard give the figure a chapter under the heading "It Costs Five Times More to Acquire a Customer than to Retain a Customer," and their conclusion is that its origin is indeterminate. What they write is that "the earliest sources that we can find attribute it to research conducted by the Technical Assistance Research Project (TARP) in Washington, D.C. in the late 1980s." Their trail runs through TARP, a 1990 Harvard Business Review article by Hart, Heskett and Sasser, and Tom Peters.

Neither Bain nor Reichheld is on it. So the assistants' answer is not a blurred memory of a real paper. It is a substitution — a famous name fitted to an untraceable number, on a page with something to sell.

Reichheld and Sasser did publish in HBR in 1990, and "Zero Defections" contains no such ratio. Its only acquisition-cost figure is absolute and industry-specific: "In credit cards, for example, companies spend an average of $51 to recruit a customer and set up the new account."

The drift starts inside the source

The claim Reichheld and Sasser did make is worth following, because it shows the corruption beginning before anyone else touches it. Their article states it three ways. The introduction: "Companies can boost profits by almost 100% by retaining just 5% more of their customers." The exhibit heading: "Reducing Defections 5% Boosts Profits 25% to 85%." The body: "Reducing defections by just 5% generated 85% more profits in one bank's branch system, 50% more in an insurance brokerage, and 30% more in an auto-service chain."

Three renderings, one article. The version that circulates today is 25–95%, which is none of them, and two of the four returned exactly that.

The laundering step downstream is public too. Amy Gallo's 2014 HBR article opens: "Depending on which study you believe, and what industry you're in, acquiring a new customer is anywhere from five

to 25 times more expensive than retaining an existing one." No footnote, no named study. Two sentences later a hyperlink appears, attached to a different claim about retention and profit. Downstream the hedge falls away and the citation slides across.

How much of this is the corpus and how much is the model

Across the wider study, nineteen of thirty-six answers carry at least one provenance error, thirty-six errors in total. Each is classified by a rule fixed before classifying: a document in that answer's own citation set either is established to contain the error, or is not. Twenty-two route to a cited document. Fourteen do not.

Those fourteen are not all fabrication. Six of the fourteen are Gemini entries, spread over four answers, and Gemini returns citation labels without URLs — so its errors cannot be traced to a document even when the same error is demonstrably inherited elsewhere. That is an artefact of one system's citation format, and it makes the split a floor rather than a measurement. A blind second coding of a quarter of the rows — by a separate language model, since no independent human reader was available — agreed on 80% of field calls, found two provenance errors this coding had missed, and is published unreconciled.

What you can do with one number

Take a figure your team quotes and ask an assistant for its original source. When it names an organisation as the source, search that organisation's own site for the number before believing it — the substitution here would have collapsed in about ninety seconds, because no Bain or Reichheld publication reporting it has been found — an absence established by search rather than proved. An attribution to an institution is not a citation. It is a name in the place where a citation should be, and it survives because opening the document is the one step nothing in the chain rewards.

The part I cannot explain

The same question was put to the same four systems in German on the following day. Claude returned four different ratios — 5, 7, 16 and 25 — and then declined to name an origin at all, saying so outright: "Leider nennen die gefundenen Artikel keine konkreten Originalstudien oder wissenschaftlichen Quellen." Unfortunately the articles found name no specific original studies or scholarly sources; the translation is mine.

The same model, one day apart, asserted a false origin three times out of three in English and refused to name one in German. Those are single runs and prove nothing on their own. But the direction is the opposite of what the corpus explanation predicts, since the German-language sources are, if anything, looser — one of them supplied a sixteenfold ratio that appears nowhere in the English material.

I do not know why. The obvious guesses — a thinner index, different retrieval behaviour, a language-conditioned difference in how readily a model will name an authority — are guesses, and the study was not built to separate them. It is the most interesting thing in the dataset and it is the thing the dataset cannot answer.

REFERENCES

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. The Quarterly Journal of Economics, 140(2), 889–942. https://academic.oup.com/qje/article/140/2/889/7990658

Gallo, A. (2014, October 29). The value of keeping the right customers. Harvard Business Review. https://hbr.org/2014/10/the-value-of-keeping-the-right-customers

Greenberg, S. A. (2009). How citation distortions create unfounded authority: Analysis of a citation network. BMJ, 339, b2680. https://doi.org/10.1136/bmj.b2680

Hart, C. W. L., Heskett, J. L., & Sasser, W. E., Jr. (1990, July–August). The profitable art of service recovery. Harvard Business Review. https://hbr.org/1990/07/the-profitable-art-of-service-recovery

Keiningham, T. L., Vavra, T. G., Aksoy, L., & Wallard, H. (2005). Loyalty myths: Hyped strategies that will put you out of business — and proven tactics that really work. John Wiley & Sons.

Liu, N. F., Zhang, T., & Liang, P. (2023). Evaluating verifiability in generative search engines. In Findings of the Association for Computational Linguistics: EMNLP 2023 (pp. 7001–7025). Association for Computational Linguistics. https://doi.org/10.18653/v1/2023.findings-emnlp.467

Peters, T. J. (1987). Thriving on chaos: Handbook for a management revolution. Alfred A. Knopf.

Reichheld, F. F., & Sasser, W. E., Jr. (1990, September–October). Zero defections: Quality comes to services. Harvard Business Review. https://hbr.org/1990/09/zero-defections-quality-comes-to-services

Stealth Agents. (2026, May 23). Customer retention cost statistics 2026. https://stealthagents.com/research/customer-retention-cost-statistics-2026

END OF ESSAY

HOW TO CITE THIS ESSAY

Isoglu, S. (2026, July 28). A vendor page supplied the source for the 5x retention rule. Sinan Isoglu. https://isoglu.com/journal/a-vendor-page-invented-the-source/

KEEP READING

From the research bench

What peer review deleted, the web kept.

https://isoglu.com/journal/what-peer-review-deleted/

From the research bench

Evidence over anecdote — what a number has to survive.

https://isoglu.com/journal/evidence-over-anecdote/

Go-to-market & pricing

Pricing is a positioning decision.

https://isoglu.com/journal/pricing-is-positioning/

Sinan Isoglu, MBA (Quantic)
Commercial growth leader, lecturer and doctoral researcher

Online version: https://isoglu.com/journal/a-vendor-page-invented-the-source/