



GO-TO-MARKET & PRICING

A flat rate buys your customer's worst month.

Flat pricing for software rested on two conditions, and inference cost broke one. The buyers who most want a flat rate are the ones with the biggest tail.

Sinan Isoglu, MBA (Quantic)

Commercial growth leader, lecturer and doctoral researcher

PUBLISHED

9 August 2026

UPDATED

20 August 2026

READING TIME

12 min

WORDS

2,718

<https://isoglu.com/insights/journal/a-flat-rate-buys-the-worst-month/>

Management summary	2
Why do pitch-deck NDR figures for usage pricing mislead buyers?	3
What conditions did practitioner models drop from the tariff-choice theorem?	3
What actually changed in consumption pricing when stated with empirical rigour?	4
How do B2B buyers respond to unpredictable billing spikes?	4
Which systematic error does a usage meter introduce into revenue forecasting?	6
Why do flat-rate bias and risk aversion identify the same target account?	7
Which two ratios should pricing leaders extract from historical billing data?	8
Boundary	8
References	9

MANAGEMENT SUMMARY

The case for charging software by the seat rather than the unit is a real result with two conditions attached: that metering costs something, and that serving more usage costs nothing. The practitioner version kept the conclusion and dropped both. Where inference cost has made the second condition false, the arithmetic changes: not because flat pricing was ever wrong, but because a flat rate is a purchase of the customer's usage variance, and it was free to hold only while the tail cost nothing to serve. The buyer side is measured: purchasing professionals pay a premium for flat rates out of an insurance motive, and what moves them is the size of the spike in their usage history, not its average. Two ratios already in the billing data decide it.

Keywords: Pricing models · Usage-based pricing · SaaS pricing · Flat-rate bias · AI pricing

The question in the room is whether to move off per-seat pricing. Someone has a slide showing that companies on usage-based pricing hold onto revenue better. Someone else points out that the finance team cannot forecast a usage line. And underneath both is an assumption nobody says out loud: that there is a right answer about which meter performs better, and that it is knowable.

For software, that question has not been answered. What has been measured is something more useful and much less discussed: what a meter actually does is decide who carries the risk when usage moves, and the party you hand it to is not indifferent about receiving it.

This is a different decision from the two next to it. Where the number sits is set by the comparison set your buyer already has in mind; moving that number up is judged before it is paid; and what your pricing page publishes signals whether buyers are segregated into visible tiers or forced onto a meter. This piece is about the unit underneath both: what you charge per.

Why do pitch-deck NDR figures for usage pricing mislead buyers?

Two figures do most of the persuading. Bessemer's State of the Cloud 2021 reports that "companies that leverage usage-based pricing typically have best-in-class net dollar retention results, averaging roughly ten percentage points higher than their traditional pricing model peer set." Bain's 2022 technology report says consumption-priced companies "increased revenue approximately 8 percentage points higher on average" than companies that mainly license or sell subscriptions.

Neither page states a sample size or a method. And read the second one closely: the companies it is describing are named in the sentence: Snowflake, JFrog, Elastic NV, Datadog, Confluent. Bain's retention figure is narrower still, holding "among the software companies with the highest valuation multiples," which selects on something the outcome is already correlated with.

These are comparisons between firms that chose different meters. Whatever separates a company that moved to consumption pricing from one that did not, whether product, market, buyer, or the decade in which it was founded, sits inside those gaps alongside the meter. That is not an accusation of bad faith; it is a description of what the quantity is, and both sources are describing their portfolios rather than running an experiment.

As far as a search conducted in August 2026 can tell, there is no peer-reviewed comparison of realized outcomes between pricing metrics in software to put next to them. Two adjacent fields do offer relevant comparisons: pharmaceutical risk-sharing agreements and usage-based motor insurance have both been studied against outcomes, but neither is part of this analysis.

So the honest position is that the horse race is unsettled. The useful part is elsewhere.

What conditions did practitioner models drop from the tariff-choice theorem?

The argument for offering a flat fee on software is not folklore. Arun Sundararajan proved it in Management Science in 2004, and the proposition is exact:

If transaction costs are nonzero, that is, if $C(q) > 0$ for $q > 0$: then it is always profit improving for the seller of an information good to offer a fixed-fee contract.

“Transaction costs” here means what it costs the seller to run usage-based pricing at all: Sundararajan’s list is monitoring and recording usage per customer, billing, payment and settlement, itemised statements, dispute resolution, the support calls that disputes generate, and the infrastructure to keep it all auditable. The result holds under weak assumptions: any positive cost function, any absolutely continuous distribution of customer types, which is why it travelled. It is a monopoly screening model under incomplete information, not a description of a competitive market, and it says offer a fixed fee, not only a fixed fee.

What travelled with it was the conclusion. What did not travel was the second condition, which sits in the model’s setup rather than in the proposition: “Variable costs of production to the firm: creating copies or providing access to the product: are zero.”

Zero. Not low, not falling: assumed away, because in 2004 for the goods he was modelling that was the right simplification.

That is not a technicality. The gain from adding a flat fee has two parts, and Sundararajan is explicit about both: you stop paying to meter the customers who take it, and you collect the surplus from the extra usage they then consume: because “the firm bears no additional variable production costs from this increase in usage.” The first part survives any cost structure. The second part is made entirely of the zero.

What actually changed in consumption pricing when stated with empirical rigour?

Sundararajan flagged the boundary himself, in a footnote: when variable costs are high relative to the cost of administering usage-based pricing, “these transaction costs affect pricing only minimally.” He was thinking of physical goods. It now describes a class of software.

Note the shape of that test. It is a ratio, not a threshold: marginal cost per unit of usage against what metering costs you to run. Most B2B software still sits comfortably on the old side of it: if your gross margin starts with a seven or an eight and your metering is mostly fixed engineering cost, you are approximately in the world the theorem describes, and the old advice is still the right advice. The products that have crossed are the ones where serving a request costs real money every time: inference-heavy features, agents that run for minutes, anything where the variable line moves with use.

If you have crossed, positive marginal cost does not repeal the result. It shrinks the second half of the gain and can turn it negative. The flat fee still saves what metering costs. What it no longer does is buy the extra usage for free.


How do B2B buyers respond to unpredictable billing spikes?

Here the evidence is unusually direct. Mario Kienzler, Christian Kowalkowski and Daniel Kindström ran four experiments with 482 purchasing professionals in Sweden and Germany: people with ten to sixteen years of purchasing experience depending on the study, buying on behalf of companies, between a third and two-thirds of which had more than 250 employees.

Offered a cloud analytics contract where the flat rate and pay-per-use came to exactly the same money, 68 percent took the flat rate. When the flat rate was made 20 percent more expensive than pay-per-use given their own stated usage, 23 percent still took it. At a 50 percent premium, 15 percent still took it. That is one experiment of 124 people split across three price conditions: roughly forty in each, so read the individual percentages as coarse

and the direction as solid. The authors put the benchmark plainly: “In the absence of biased choice, none of these purchasing professionals would be expected to choose the flat rate in the two premium conditions.”

Figure 1 The preference at equal money, and what a premium does to it




Kienzler, Kowalkowski & Kindström (2021), *Journal of Business Research* 132, Study 1, p. 407. Scenario-based choices by purchasing professionals, $n = 124$, about forty per condition. Single shares read as direction: a presentation-order effect elsewhere in the paper moved them between 45% and 73% (fn. 9).

Two things to keep hold of before that becomes a licence to raise prices. First, price bites: 68 to 23 to 15 is a collapse, not a plateau, and it is what makes the preference something you can price rather than something you can exploit indefinitely. Second, all of this is scenario-based: hypothetical purchases, self-reported choices, no money at stake, and the authors ask for field validation. A presentation-order effect in one of the studies moved the share between 45 and 73 percent depending on which option was shown first, which is roughly how precisely any single percentage here should be read.

The finding that matters for a pricing decision is the second experiment. Every participant saw the same average past usage: 500 hours a month. What varied was the shape around it: one group saw a range of 100 to 600 hours, the other 400 to 900. Same mean, bigger spike. With the bigger spike, 89 percent chose the flat rate at equal cost: 37 people in that cell. Without it, 56 percent, which is not statistically distinguishable from a coin flip. The difference is tested across the whole sample of 145 rather than eyeballed between two cells, and it holds in the regression.

Figure 2 Same average, different tail



Kienzler, Kowalkowski & Kindström (2021), Study 2, p. 408; $n = 145$, scenario-based. Every participant saw the same 500-hour average; only the bounds varied. The contrast is tested across the whole sample: $\chi^2(1, N = 145) = 9.51, p < .01$; logistic $B = 1.19$, 95% CI 0.39 to 1.98.

That second experiment also measured three candidate explanations for the preference: insurance, convenience, and dislike of the running meter. Only one was significantly related to choosing the flat rate across all four of its conditions: the insurance effect.

So the customers reaching for a flat rate are not necessarily your heavy users. They are the ones whose usage has the largest upside relative to its average, and what they are buying is not simplicity. It is cover.

Which systematic error does a usage meter introduce into revenue forecasting?


The obvious worry about buyers who overpay is that they eventually notice and leave. Anja Lambrecht and Bernd Skiera checked the actual transaction records of an internet provider, and the answer is the opposite of intuition.

More than half of the consumers with a flat-rate bias paid at least 100 percent more than they would have on the cheapest tariff available to them. They did not churn more than anyone else. Their customer lifetime value ran 87 to 135 percent higher than customers who picked the cheapest tariff. The damage sat entirely on the other side of the error. Consumers who wrongly held pay-per-use, even though a flat rate would have saved them money, churned at monthly rates 340 to 1,040 percent higher, and their lifetime value came out slightly negative. A nested logit puts the two responses in different places: the mispriced flat-rate customer, if they act at all, switches plan with the same provider; the mispriced usage customer leaves.

Two cautions, because this is the number most easily misread. “Pay-per-use bias” means choosing pay-per-use when a flat rate would have been cheaper. It is a customer’s plan-choice error, not a property of usage pricing; the finding is not that metering causes churn. And the lifetime-value damage is the least robust figure in the study: across the authors’ sensitivity range it straddles zero on both of their criteria, running -13 to $+5$ percent on one and -8 to $+14$ on the other, while the flat-rate gain stays large throughout. This is consumer internet access from the early 2000s, and it is one provider.

What it establishes is the thing the flat-rate advice was always quietly right about. Under a zero-marginal-cost product, a customer who overpays on a flat rate is not a leak. They are the most valuable customer you have, and they stay.

Figure 3 The two plan-choice errors are not the same size



Lambrecht & Skiera (2006), Journal of Marketing Research 43(2), p. 215, version of record (a circulating manuscript reads 46.6%; the typeset article reads 46.4%). One internet provider, consumer access, early-2000s transaction records.

Why do flat-rate bias and risk aversion identify the same target account?

Put them together and the mechanism is not complicated.

A flat rate is a purchase. You are buying your customer's usage variance from them, and they are measurably willing to pay for the transfer. The buyers most willing to pay are the ones with the biggest spike relative to their average. And absorbing that spike cost you nothing at all: for exactly as long as serving a unit of usage was free.

That was the whole trade. It was a good one. It is still a good one wherever the second condition holds.

Where it no longer holds, the same measured pull now sorts your customers in a direction you would not have chosen: the accounts that select into your flat tier are disproportionately the ones whose worst month you now have to pay for in cash. That is not an argument against offering a flat tier: Sundararajan's result still says you should, and Lambrecht and Skiera say those customers will be loyal. It is an argument that the flat component of your menu has to be priced against the tail rather than against the average, and that the price has to be a real one, because 15 percent will pay a 50 percent premium and 68 percent will not.

There is a version of this that goes wrong, and it is worth naming: a hybrid menu priced off average usage, sold to the customers who chose it for its insurance value. That is not a pricing model. It is an unpriced option, written by you, exercised by whoever has the most volatile demand.

Which two ratios should pricing leaders extract from historical billing data?

Neither of the numbers this argument turns on is a number you have to go and find.

Peak-to-average usage, per account. Take each customer’s busiest month against their own average: the ratio Kienzler’s experiment actually manipulated, 600 over 500 in one condition and 900 over 500 in the other. Sort your accounts by it. That column tells you which customers your flat tier is selecting for, and of everything about a customer that might drive the preference, it is the one property anyone has shown to move it. The other lever that moves it is the premium you charge, and that one is yours.

Marginal cost per unit of usage, against what metering costs you to run. Not gross margin: the cost of one more unit, against the annual cost of the metering, rating, billing and dispute machinery you would need to charge for it. If that ratio is small, you are in the world the theorem describes and this essay is about somebody else. If it is not, you already know which of your customers you cannot afford to insure at the average.

The inputs were never missing. What was missing was the condition that made ignoring them safe.

Figure 4 Two ratios, read off your own billing data

ACCOUNT OR SEGMENT	PEAK-TO-AVERAGE USAGE	MARGINAL COST PER UNIT	METERING RUN-COST
Your largest first.	Busiest month over own average: 600/500 vs 900/500 in the experiment.	What one more unit of usage costs you to serve.	Rating, billing and dispute machinery, per year.

The first ratio is the one customer property shown to move the preference; the second decides whether absorbing the variance is still nearly free. Price the flat component against the tail, not the average.

Author's own worksheet.

Boundary

Boundary. The buyer preference for flat pricing is measured; the correct meter for a given product is not. Calculate the two billing ratios on your own usage distribution and test the tail before changing the model.

Evidence base. The analytical frame also draws on these additional sources: Sundararajan 2004. The links identify the exact works; they support the mechanisms and boundary conditions discussed here, not every claim in isolation.

REFERENCES

- Bain & Company. (2022). Is consumption-based pricing right for your software? Technology Report 2022. Page capture 2026-08-08 via the Internet Archive; the live page returned HTTP 403. <https://www.bain.com/insights/is-consumption-based-pricing-right-for-your-software-tech-report-2022>
- Bessemer Venture Partners. (2021). State of the cloud 2021. BVP Atlas. Page capture 2026-08-08. <https://www.bvp.com/atlas/state-of-the-cloud-2021>
- Kienzler, M., Kowalkowski, C., & Kindström, D. (2021). Purchasing professionals and the flat-rate bias: Effects of price premiums, past usage, and relational ties on price plan choice. *Journal of Business Research*, 132, 403–415. <https://doi.org/10.1016/j.jbusres.2021.04.024>
- Lambrecht, A., & Skiera, B. (2006). Paying too much and being happy about it: Existence, causes, and consequences of tariff-choice biases. *Journal of Marketing Research*, 43(2), 212–223. <https://doi.org/10.1509/jmkr.43.2.212>
- Sundararajan, A. (2004). Nonlinear pricing of information goods. *Management Science*, 50(12), 1660–1673. <https://doi.org/10.1287/mnsc.1040.0291>

END OF ESSAY

HOW TO CITE THIS ESSAY

Isoglu, S. (2026, August 9). A flat rate buys your customer's worst month. Sinan Isoglu.
<https://isoglu.com/insights/journal/a-flat-rate-buys-the-worst-month/>

KEEP READING

Go-to-market & pricing

The integration savings come from what you put in

<https://isoglu.com/insights/journal/integration-savings-come-from-what-you-put-in/>

Go-to-market & pricing

What is usage-based pricing?

<https://isoglu.com/insights/journal/what-is-usage-based-pricing/>

Go-to-market & pricing

What is value-based pricing?

<https://isoglu.com/insights/journal/what-is-value-based-pricing/>



Sinan Isoglu, MBA (Quantic)

Commercial growth leader, lecturer and doctoral researcher

© 2026 Sinan Isoglu · isoglu.com · All rights reserved

Online version: <https://isoglu.com/insights/journal/a-flat-rate-buys-the-worst-month/>