



Does Customer Maturity Cause Renewal?

Replacing a company-wide belief with a number the business could plan around — identification design, five estimators, and the path from estimate to executive adoption.

+12.9pp

causal renewal lift

5

estimation methods

\$323M

at-risk ARR studied

11,166

renewal events



THE SETTING

01

A \$300M renewal book, riding on an unproven thesis

THE POINT

The job was not to bless the thesis or kill it — it was to find out which parts of it deserve real money.

The operating thesis. Customer-success strategy — and a major program investment case — rested on the belief that moving customers up a maturity model protects renewal. Everyone could see the correlation; no one had shown causation.

The stakes. Mature customers renew better, but they are also bigger, better supported, and healthier to begin with. If maturity is only a marker of healthy accounts, the program funds motion that would happen anyway.

My role. I owned the question end-to-end: translated the thesis into an estimable causal question, designed the identification strategy, ran five estimators, and carried the result into renewal planning.

“

If we cause a customer's maturity to advance, does their renewal probability actually move — and by how much, where, and for whom?

11,166

contract-renewal events

7,200

customer accounts

\$323M

at-risk ARR in cohort

~20 mo

of renewal outcomes



One clock per contract, read before the decision

THE POINT

The design decisions do the causal work before any model runs — reverse causation is blocked by construction, not assumption.

Treatment: a 0.5+ advance on the composite maturity score, in a window closing three months before renewal.



01

Movement, not level

High-maturity accounts renew better — but they are also different accounts. Comparing accounts that moved against comparable flat accounts isolates a lever the business can pull, not a trait it can admire.

02

Why the 3-month lag

A customer heading into renewal is already engaged, which can masquerade as a maturity gain. Reading treatment before the renewal conversation blocks reverse causation by design.

03

Who is compared

Treated: advanced 0.5+ (1,156 events). Control: stayed flat (10,010). Decliners are excluded — a different process; lumping them into control would manufacture a larger effect than is real.



DESCRIPTIVE FIRST


03

The raw gap, before any modeling

THE POINT

Pooled raw gap: +7.4pp. Honest — but the groups were not alike to begin with, and the raw number turns out to understate the effect.

Renewal rate, advanced vs. stayed-flat, by starting maturity band — 11,166 events.



START 0-1

+20.6pp

659 advanced · 3,239 flat

START 1-2

+3.1pp

426 advanced · 5,113 flat

START 2-3

-1.1pp · thin

71 advanced · 1,658 flat




A fairer comparison makes the effect bigger

THE POINT

Diagnosing the direction of selection — not just its presence — changed the planning number by nearly a factor of two.

Adjusting for confounders pushed the estimate up. Advancers started at lower maturity — and low-maturity accounts renew worse — so the raw comparison pitted them against a healthier flat group and hid the effect.



Why this read matters. The naive story — “healthy accounts advance, so the raw gap overstates” — is exactly backwards in this data. Getting the direction of selection right is the difference between under-selling a real lever and over-selling a fake one.



TRIANGULATION


05

Five methods, one result

THE POINT

A method resting on entirely different assumptions still clears zero comfortably — the strongest answer to “isn’t this just healthy accounts?”

Run independently and asked to agree. Cluster-bootstrap 95% intervals, resampled by account.





STRESS TESTS

06

Trying to break the result, on purpose

THE POINT

Robustness here is not a checklist — each test was chosen to give the confounding objection its best shot.

TEST 1 · DELETE THE STRONGEST CONFOUNDER

+12.9 → **+11.2** *pp*

Baseline maturity is the biggest measured driver of both advancing and renewing. Deleted entirely, the re-run full pipeline moves 1.7pp — the design and remaining covariates recover nearly the same effect. This is not baseline maturity in disguise.

TEST 2 · WHERE HUMAN TOUCH CANNOT REACH

A champion or attentive CSM story predicts the effect should vanish where there is little of either. It doesn't:

Enterprise · high-touch	+12.0 pp
Mid-Market	+7.0 pp
Scaled · pooled CS	+14.0 pp
Self-Serve · no CSM	+15.6 pp

The honest residual. Customer-side budget health is only partly proxied by size and ARR — named openly rather than buried. A factor strong enough to nullify the effect would need to out-predict baseline maturity on both legs at once while staying uncorrelated with the model. Unusual, and a finding in itself if it existed.

VERDICT *The estimate survives deletion of its strongest confounder and holds where the human-touch story cannot reach.*



WHERE IT HOLDS

07

The average is true — but it is not the targeting answer

THE POINT

One pooled number would have pointed spend at \$200M+ of flat ARR with a near-zero effect. Decomposition turned a true statistic into a correct allocation.

Pooled +12.9pp is decision-grade for planning. For targeting, the effect concentrates where accounts still have room to move.

STARTING 0–1

+20.9pp

PRIMARY TARGET

Raw (+20.6) and adjusted (+20.9) agree, methods converge, and the flat group renews at only 63% — large headroom.

STARTING 1–2

≈ 0 pooled

TARGET THE SUBSET

Accounts already renew ~81–84%. But the average hides a longer-tenured, higher-value subset (n≈1,050) with a real +9.1pp adjusted lift. That subset — not the whole band — is the secondary target.

STARTING 2–3

n = 71

OPEN QUESTION

Too thin to estimate — not a negative finding. Advancing already-mature accounts should be justified on adoption or expansion grounds, not renewal.



08

WHAT IT IS WORTH

The problem, the lever, and the prize

THE POINT

Take-up is the one unmeasured input — sized next via CS goals, recent advancement rates, and randomized nudge experiments.

Effect, base, and dollars all come from the same study, population, and unit — no cross-source projection.

THE PROBLEM

\$290M

at-risk ARR sitting flat — the dollars a program could still move

THE LEVER

~13¢/\$1

renewal protected per dollar of ARR advanced, per cycle

THE PRIZE

\$9–28M

protected renewal per cycle, depending on how many accounts a program moves

PROTECTED RENEWAL PER CYCLE, BY TAKE-UP OF THE FLAT

¼ of accounts move



~\$9M

half move



~\$19M

¾ move



~\$28M

ALREADY REALIZED

\$4.3M

protected ARR on the book that advanced — +12.9pp
× \$33M advancing ARR. Counterfactual renewal
71.8% → actual 84.4%.

ONE ASSUMPTION The measured lift applies to comparable not-yet-advanced accounts. Everything else in the sizing is measured, same-population, same-unit.



FROM ANALYSIS TO ADOPTION

How the result is being used

THE POINT

Analysis that doesn't change a decision is a draft. This one now sits inside the plans that fund the work.

A causal estimate only matters if the organization plans on it. The work was built to travel — and it did.

01

Embedded in the executive churn review

Finance pulled the segment-level causal results directly into the company's executive churn analysis — maturity advancement now appears alongside coverage and adoption as the one causally proven retention lever, with tier-level lifts and protected-ARR figures cited as decision inputs.

02

Feeding a board-level action plan

The CFO's office is using the findings to shape the churn action plan presented to the board — the causal lift underpins the retention pillar of that plan, exactly the bar this analysis was designed to clear.

03

Adopted as the working planning value

The pooled lift was adopted as the standing retention value of a 0.5+ maturity advance — a stable number renewal forecasting reuses rather than re-derives, with the band-level decomposition governing where plays are targeted.

04

Experiment pipeline designed

Randomized nudge experiments and CS advancement-rate baselining are scoped to pin down take-up — converting the \$9–28M range into a defensible planning number and closing the loop from “advancement causes renewal” to “our intervention causes advancement.”



What this analysis does not claim

THE POINT

Keeping the causal boundary explicit is what made the number trustworthy enough for finance and the board plan to build on.

Knowing exactly where the evidence stops is what lets us stand behind it where it holds. Each is a scoped follow-on, not a caveat that weakens the headline.

§1

Advancement → renewal, not intervention → advancement

The study proves moving maturity moves renewal. It does not yet prove a specific CS or product play causes maturity to move — that is what the randomized nudge experiments are for.

§2

Decline is a separate process

Declining accounts renew ~8pp below flat descriptively, but were deliberately excluded, not causally estimated. Losing maturity deserves its own design, not a free ride on this one.

§3

Level-as-protective is unanswered

High-maturity accounts renew well, but they are also bigger and better supported. This study measures movement; whether sitting at high maturity protects renewal is a different, confounded question.

§4

Dose-response is out of reach, for now

A 0.5 advance is treated as one uniform move. Whether the payoff saturates at higher maturity is an effect-shape question this binary design cannot see — designed, awaiting sample.

01

Turn beliefs into estimands

The highest-leverage act wasn't a model — it was converting “maturity drives retention” into a precisely defined treatment, window, comparison group, and estimand the whole company could argue with.

02

Triangulate, don't decorate

Five estimators with different failure modes, run independently and asked to agree, plus deliberate attempts to break the result. Convergence is what made one number decision-grade.

03

The average is not the allocation

Pooled +12.9pp is true and plannable; the targeting answer lives in the decomposition — separating those two uses prevented spending against \$200M+ of near-zero-effect ARR.

04

Adoption is the deliverable

The result now sits inside the executive churn review and a board-level action plan, with a named path from measured effect to funded program.

“

The job was not to bless the thesis or kill it — it was to replace a belief with a number the business could plan around, and to be exact about where that number stops.