PORTFOLIO CASE STUDY · PART II · HETEROGENEITY & TARGETING

Where the Effect Actually Lives

Part I proved advancing maturity causes renewal. Part II asks the harder question — for which accounts, and which dollars? A causal tree on doubly-robust scores, and the targeting map it produced.

+25.5pp
strongest leaf effect

4 leaves
one causal tree

82%
of flat $ where effect is weak

$10–16M
grounded opportunity / cycle

William Gieng · Data Science & Analytics · B2B SaaS

All data shown is synthetic — simulated to preserve the original analytical patterns; no proprietary figures, names, or segments are disclosed.

PICKING UP FROM PART I

01

One proven number — now, who exactly?

THE POINT

Part I answered “does it work?” Part II answers “on whom?” — because those are different questions with different owners.

Part I established a pooled +12.9pp causal renewal lift from a 0.5+ maturity advance, robust across five estimators. Averages fund programs; they don’t aim them. The effect concentrates where accounts still have headroom:

STARTING 0–1	STARTING 1–2	STARTING 2–3
+20.9pp	**≈ 0 pooled**	**n = 71**
63% flat-group renewal floor — large headroom, and raw and adjusted estimates agree. This is where the lever bites.	Accounts already renew ~81–84%. The average hides a longer-tenured, higher-value subset with a real +9.1pp lift — pooled bands blur what a tree can separate.	Too thin to estimate. Advancing already-mature accounts is an adoption or expansion case, not a renewal one.

Why bands aren’t enough. Bands are one-dimensional. Effects in this data vary by starting maturity and contract value and tenure at once — so Part II lets the data draw its own boundaries: a causal tree partitions all 11,166 contract-renewal events by estimated impact, not by outcome or by intuition.

THE METHOD

02

A causal tree, grown on doubly-robust scores

Each event gets a doubly-robust score — its contribution to the causal estimate. A shallow tree splits on interpretable pre-renewal traits (starting maturity, contract value, tenure) to find groups whose estimated lift genuinely differs.

THE POINT

Every contract lands in exactly one leaf — grouped by estimated causal impact, using traits measured before the renewal conversation.

THE LEAVES

03

Four segments, four different stories

THE POINT

Effect size, dollars, and confidence never line up in the same place — which is exactly why the tree beats a single number.

Leaf A — primary target

+25.5pp [21.4, 29.7]

Very low starting maturity, smaller contracts. 401 advancers, $15.9M flat ARR at a 59% renewal floor. The clearest, largest effect — and the cheapest place to be wrong.

Leaf B — primary, higher value

+15.7pp [9.4, 23.5]

Low maturity with real contract value: $38.2M flat ARR at a 65% floor. Wider interval — cite the range, not a point — but strong and screened.

Leaf C — modest tail, screen it

+4.9pp [2.2, 8.0]

The mass of the book: 7,377 events, $129.6M flat at an 82% floor. Real but small per contract; the long-tenure, >$25k subset ($41.2M flat) concentrates the effect at ~+9pp — target the subset, not the leaf.

Leaf D — do not target

−3.1pp [−12.8, +4.3]

The biggest contracts, already renewing at 86%. No detectable lift — a maturity-led renewal play here funds motion that would happen anyway. Adoption and expansion are the right lenses.

04

THE INCONVENIENT MAP

The biggest pools sit where the effect is weakest

$290M of flat ARR, partitioned by leaf. The strong-effect leaves hold less than a fifth of the money.

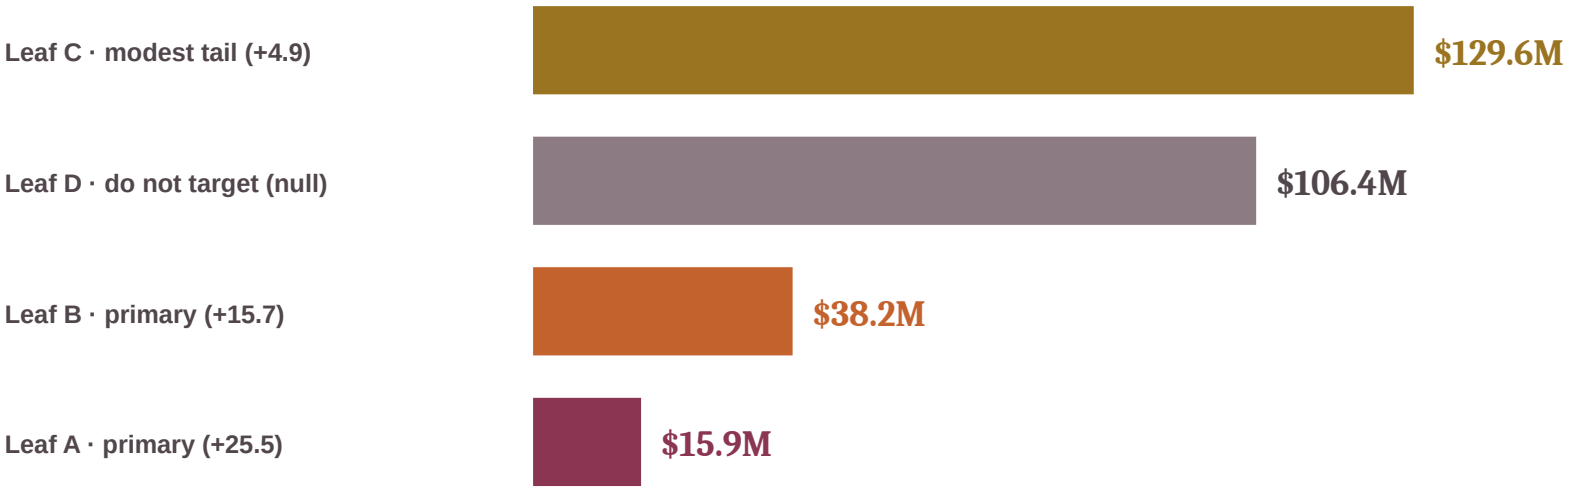

82%

of flat ARR ($236M) sits in Leaves C and D, where the per-contract effect is modest or absent — the pool a naive rollout would spend against.

19%

of flat ARR ($54.1M) sits in Leaves A and B — screened, strong, trim-stable effects. Small pool, high certainty: where a renewal-led play earns its keep.

THE POINT

If the dollars and the effect lived in the same place, this would be an easy program. They don't — and saying so is the analysis' job.

SIZING, HONESTLY 05

How big is the opportunity, really?

FULL FLAT BOOK

$290M

every currently-flat contract, all tiers and products — the number that sounds best and means least

→

EXPECTED CHURN DOLLARS

~$56M

what the model expects to be lost from the flat book absent any intervention — the real size of the problem

→

CAUSALLY GROUNDED

$10–16M

leaf effect × leaf flat ARR: $10.1M in Leaves A+B alone; $16.4M adding Leaf C's modest tail — per cycle, at full reach, before take-up

How this reconciles with Part I's $9–28M. Part I scaled one pooled effect by take-up scenarios on the whole flat base. Part II replaces that with leaf-level effects on leaf-level dollars — a tighter, better-aimed range that lands inside Part I's envelope. Same study, sharper aim.

THE POINT

Neither the full flat base nor total at-risk dollars is the right number. The defensible figure is effect × dollars, in screened leaves only.

READ IT RIGHT *$10–16M against $56M of expected churn is the honest denominator — not the $290M that mostly renews regardless. Compare the lever to doing nothing, not to the whole base.*

TARGETING, TODAY 06

Effect size alone does not pick the pilot

THE POINT

Accounts likely to advance on their own are poor targets even with a strong effect — incrementality, not effect size, picks the pilot.

A movability screen — modeled advancement propensity among flat lines — splits each leaf's dollars by whether accounts would likely move without help:

Assisted pilot · A+B stuck movers — **$5.2M** — PILOT NOW

Strong-effect leaves, low organic-movement signal — the accounts that gain a lot but won't get there alone. Cleanest incrementality; smallest, sharpest test.

Incrementality holdout · A+B organic movers — **$48.8M** — RANDOMIZE

Most strong-leaf dollars sit with accounts already likely to advance. Fund a randomized holdout — measure what the program adds over organic movement before scaling spend here.

Screened next wave · Leaf C subset — **$41.2M** — SECOND WAVE

Long-tenure, >$25k contracts inside Leaf C, where the modest leaf effect concentrates at ~+9pp. Queued behind the pilot, not alongside it.

Excluded · Leaf D and the C remainder — **$195M** — DO NOT FUND

Null effect or thin per-contract value at an 82–86% renewal floor. Excluded by rule in the pipeline — not by a filter applied after the fact.

07

THE CAUSAL BOUNDARY

What is established — and what still needs testing

THE POINT

The study proves advancement drives renewal in specific leaves. It does not yet prove any given play drives advancement — keep that line bright.

ESTABLISHED

- Low-maturity advancement (Leaves A and B) carries the clearest, strongest renewal lift — +15.7 to +25.5pp, screened and trim-stable.
- The renewal value is not evenly distributed: dollars concentrate where effects don't, and the evidence supports concentrating renewal effort at the bottom of the maturity scale.
- Leaf D — the largest contracts — shows no detectable lift and should not be prioritized for a maturity-led renewal play.

STILL NEEDS TESTING

- Whether a specific CS or product action causes an account to advance — the study covers advancement → renewal, not intervention → advancement.
- Whether the movability screen is right where being wrong is most expensive — validated by adding organic-mover arms to the randomized pilot.
- Whether the leaf boundaries hold outside the build sample — a held-out or randomized test, before the targeting rules become policy.

NEXT *A randomized pilot on the assisted pool with organic-mover holdout arms — the experiment that converts this map into funded policy.*

IN CLOSING · PART II

01

Let the data draw the boundaries

Bands were the human guess; the causal tree found the real structure — maturity × contract value — and every contract now lands in exactly one leaf with its own evidence.

02

Dollars concentrate where effects don’t

82% of flat ARR sits in weak- or no-effect leaves. Naming that is what separates a targeting map from a sales pitch — and protects $195M from a play that wouldn’t pay.

03

Incrementality picks the pilot

The strongest-effect accounts are also the likeliest to advance on their own. Effect size funds the program; the movability screen and a randomized holdout decide where it starts.

04

Small, sharp, then scale

$5.2M assisted pilot → $41.2M screened second wave → policy. Each step gated by an experiment, not by enthusiasm.

> *Part I earned the right to spend. Part II earned the right to aim — and drew the line where spending stops.*