

What your churn data knew, and when it knew it

Prepared for the Board • Period 1 Jan 2025 to 1 Aug 2026

Accounts analysed 600 **Churn events** 120 **Method** Walk-forward, survivorship-safe

Issued 9 Aug 2026

Fixture, not a client report. Every figure below is computed from seeded demonstration data (seeded demonstration tenant, card on file, monthly) and describes no real company. It exists to show the structure of the deliverable and to exercise the sections that a thin dataset cannot fill.

1 VERDICT

Over the period analysed you lost \$298K of ARR to churn, a gross revenue retention of 88.8%. Of that, 17% was already flaggable 90 days before cancellation, against a random baseline of 6%.

\$298K	88.8%	17%	40d
ARR lost to churn	Gross revenue retention	Churned ARR flaggable at 90 days	Median first intervention before churn

2 DATA SUFFICIENCY

Conclusions in this report are drawn only where the underlying data supports them. Cells below the reportability floor are marked and left unquantified rather than printed with confidence intervals too wide to act on.

No per-source coverage map was produced for this period, so the table of sources is omitted rather than filled with assumptions. The reportability floor below still applies to every figure in this report.

REPORTABILITY FLOOR

A cell is printed only when it rests on at least **20 events**. Below that the cell is named and left unquantified. This is why some sections of this report carry a reason where a reader might expect a number: the figure was not withheld, it was never supportable.

3 WHAT YOU LOST

Churn is not one phenomenon. Voluntary cancellation, payment failure, and downgrade have different causes and different fixes, and conflating them produces retention programmes aimed at the wrong third of the problem.

SHARE OF CHURNED ARR



SHARE OF TOTAL ARR LOST, A DIFFERENT DENOMINATOR




ARR lost by category. The percentages above the rule are shares of churned ARR. Contraction sits below the rule because it is a share of total ARR lost: it happens to accounts that did not churn, so it does not belong in the same total. Categories that could not be measured are listed below and are excluded from both the bars and the shares.

Every churned account had a resolvable ARR value, so nothing fell into an unclassified bucket.

Gross revenue retention: **88.8%**, on a starting book of \$5.09M.

4 WHAT WAS KNOWABLE

For each account that churned, the model was rebuilt using only information that existed at the time, with no later data and no survivorship. The figures below show how much of the churned ARR was already in the flagged set at each horizon, beside what a randomly chosen set of the same size would have captured.



Share of churned ARR present in the flagged set at each horizon, against a random set of identical size. The baseline is drawn beside each bar rather than in an appendix: a lead-time figure without its control is not a finding.

HORIZON	FLAGGED SHARE	RANDOM BASELINE	EVENTS
30d	64%	15%	90
60d	35%	10%	60
90d	17%	6%	35

5 HOW FAST YOU REACTED

Detection only matters if it is followed by contact. Against the accounts that churned in this period, this is what the record shows.

40d

Median days from first intervention to cancellation

\$218K

Churned ARR with no recorded intervention

73%

Share of churned ARR unattended

READ THIS BEFORE ACTING ON THE NUMBER ABOVE

This is a measure of *recorded* intervention. Conversations held over email, phone or in person without a corresponding record are invisible here and will inflate the "no intervention" figure. Treat the unattended ARR as an upper bound, not a settled fact.

6 MODEL TRACK RECORD

8 walk-forward folds, each trained only on data preceding its origin date, with an embargo between the training and scoring windows to prevent boundary leakage.

METRIC	RESULT	INTERPRETATION
Precision @ k	44%	95% CI 38%–50%
Lift over base rate	3.1×	Base rate 14%
ROC AUC	0.71	Across all valid folds
Separable from chance	Yes	Interval excludes the baseline
Negative control	0.497	Shuffled labels, no residual signal
Leakage canary	Clean	No feature exceeds its as-of availability

The negative control and the leakage canary are reported beside the accuracy figures rather than in an appendix, because their failure would invalidate every row above them.

7 WHAT THIS DOES NOT PROVE

This report establishes what your data could have revealed and when, not what would have happened had you acted. Accounts that received an intervention are not comparable to accounts that did not: your team chose which to work, and that choice almost certainly tracked how salvageable the account already looked. Every figure here is therefore correlational. Measuring the effect of an intervention requires withholding it from a randomly chosen group, which this analysis does not do and which no retrospective analysis can do after the fact.

Specifically:

- **The reaction-time findings are correlational.** Accounts that received early contact may differ systematically from those that did not: larger, better covered, more visible. That selection effect is not controlled here.
- **Flagged-and-saved is not evidence of a save.** Some flagged accounts would have renewed regardless. Without a control arm, that fraction is unknown.
- **Historical accuracy is not a forward guarantee.** It measures a model against a period that has already happened, under conditions that may not persist.

Establishing causation requires a randomised design running forward: a stratified share of flagged accounts assigned to a different priority before outcomes are known. That is the only design that answers whether the intervention worked, and this report is not a substitute for it.

8 SEALED FORWARD FORECAST

The commitment below was sealed on 9 Aug 2026, before its outcome was known. It will be scored on 7 Nov 2026, 90 days from the seal, so the next edition of this report can be checked against this one rather than quietly replacing it.

COMMITMENT	VALUE	BASIS
Accounts in the flagged set today	214	Sealed as a hashed list, not by name
ARR they represent	\$1.18M	Exposure, not a prediction of loss
Expected to churn within 90 days	81 to 107	This backtest's own measured precision interval
A random set of the same size would yield	30	Base rate 14%
Result published regardless of sign	Yes	See below

Sealed account list (SHA-256):

```
ebddd0896887876e87fd371d075b39e02213c7d322d7c172f89c06893a7bf77e
```

Chain head:

```
37a46b32eab285e8b7c521e43ea0f35e1e67fdb2600053b505800e5a4f10b3d3
```

THE RULE THIS COMMITS US TO

This result is reported on the date above whatever its sign, including if it is worse than chance. Re-hashing the delivered account list must reproduce the value above; a mismatch means the list was altered after sealing and the result should be discarded. The list itself is not printed here, so this page can be circulated without carrying account names.

Method. Account state was reconstructed to its as-of value at each prediction date from an append-only event ledger. Labels were derived from billing and subscription events, never from present-day account status. Evaluation is walk-forward with out-of-time splits, an embargo between training and scoring windows, and calibration fitted within each training window only.

Fixture · seeded demonstration data (seeded demonstration tenant, card on file, monthly) · no real company · Prepared 9 Aug 2026 · Indexxero