Turning Subscriber Churn into Retention Strategies

Subscriber Insights and Predictive Modeling for Streaming Retention

William Gieng | Data Science and Analytics

All data shown is synthetic and anonymized to preserve confidentiality.

Agenda

OVERVIEW

01 Project Overview

- Subscriber churn prediction overview
- Objective
- Available data

02 Subscriber Insights and Analysis

- Subscriber segmentation
- Engagement analysis
- Acquisition and billing insights
- Subscriber survival curves

03 Predictive Modeling

- Predictive model results
- Implications and next steps

01

Project Overview

Context, objective, and available data

Why Subscriber Churn Matters

CONTEXT

- Today's streaming customers can cancel with ease — competition is only a click away
- Reducing churn directly boosts lifetime value (LTV) and lowers costly acquisition needs
- Identifying at-risk subscribers early lets the business deliver targeted offers that keep them engaged
- Proactive retention strategies are essential — profitability depends on it
- Engagement and viewing trends provide early warning signals of potential churn

KEY QUESTION

Can we spot these churn signals early enough to act on them?

The answer drives everything that follows.

Objective: Predict, Understand, and Act

OBJECTIVE

01

Define the Problem

Subscribers are leaving (churning) too soon.

02

Explore Solutions

Can we use data to understand which subscribers are likely to churn?

03

Evaluate Impact

What actions (offers, engagement, personalization) would improve retention while reducing costs?

04

Take Action

Build predictive models to identify at-risk subscribers and execute proactive retention strategies.

Available Data: Scope and Limitations

DATA

~30K

subscriber records

Q1 2023 cohort

4

data domains

engagement, billing, demographics, lifecycle

~3.5 mo

observation window

limits seasonal trend detection

Subscriber Engagement

Hours watched, distinct shows and channels, DVR saves, stream starts. Strong indicators of retention vs. churn.

Payments and Lifecycle

Payments, refunds, billing method and platform, acquisition method and platform. Indicators of loyalty and cancellation behavior.

Demographics

Limited or unavailable. Access to household and location data could improve churn profiling and predictive power.

Historical Depth

More history would reveal seasonal trends and strengthen predictive power across full subscription lifecycles.

02

Insights and Findings

Segmentation, engagement, acquisition, and survival analysis

Subscriber Segmentation: Defining the Cohorts

SEGMENTATION

PRIMARY FOCUS

Rapid Churners

Subscribers who cancel after just 1 month.

PRIMARY FOCUS

Retained Subscribers

Subscribers who stay for 3+ months (into month 4 and beyond).

SECONDARY

Moderate Churners

Subscribers who churned after 2 months.

EXCLUDED

Refunded / Trial Only

Never converted after trial period; excluded from modeling.

Note: Insights analysis covers all groups except Refunded/Trial Only. Predictive modeling focuses on Rapid Churners vs. Retained Subscribers.

Engagement

Not all engagement equals loyalty.

Not All Engagement Equals Loyalty

ENGAGEMENT

High early usage may still lead to churn if it reflects binge-watching behavior.

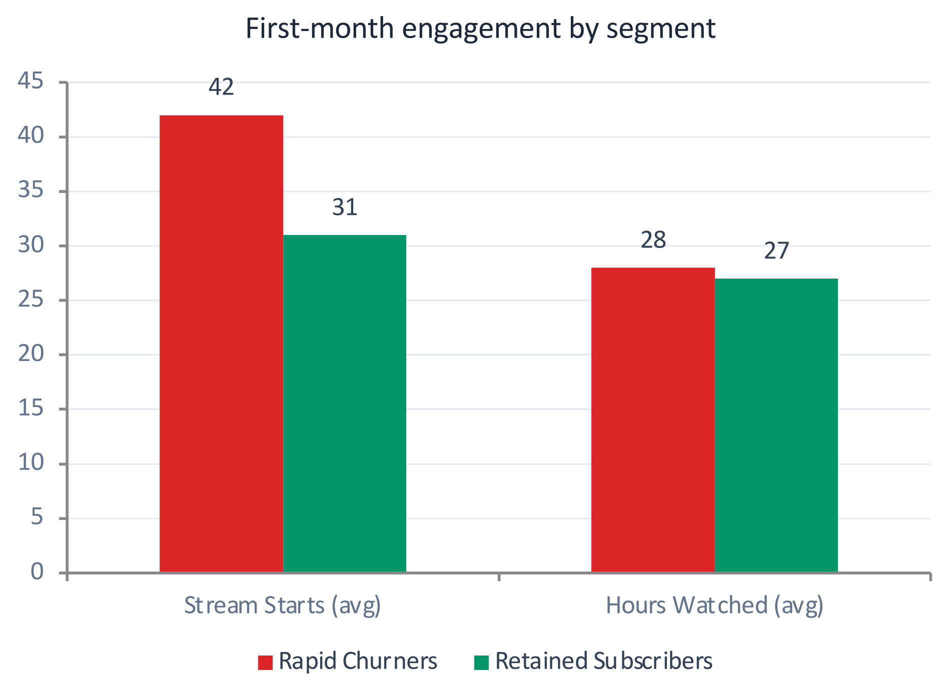

Rapid Churners start strong

More stream starts and slightly higher hours watched in first month vs. other groups.

Possible behavior

Binge-watching favorite shows during first month, then canceling.

Implication

Retention risk isn't just about low engagement — high-engagement users can churn once they've consumed what they came for.

Takeaway: First-month engagement patterns reveal who is testing the service short-term vs. who is building lasting habits.

Engagement metrics normalized by active months. Synthetic values for illustration.

Engagement Distribution Insights

ENGAGEMENT

Wide spread of usage patterns shows that churn risk is linked to behavior, not just total activity.

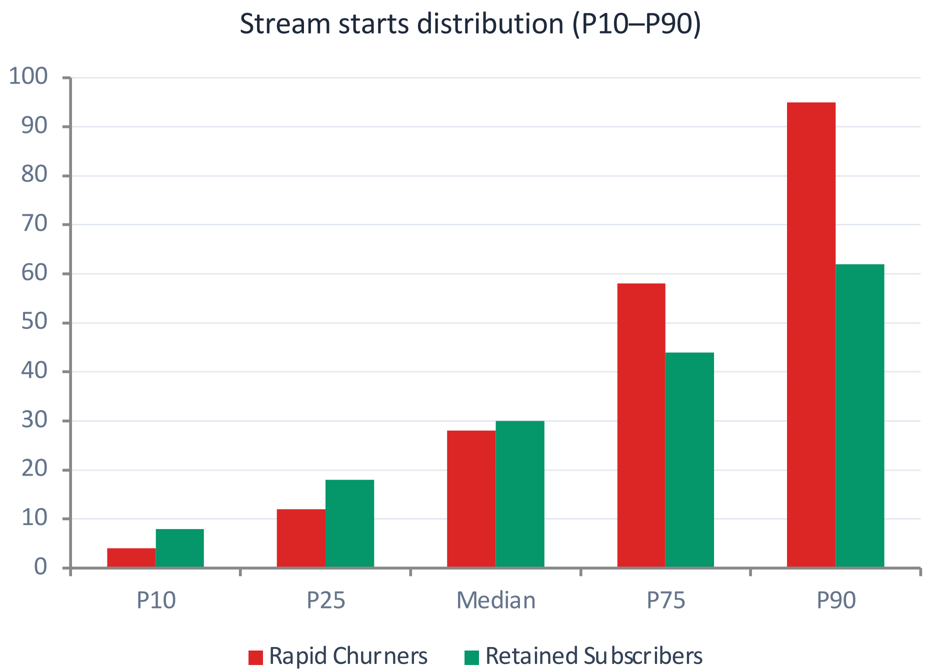

Stream Starts

Rapid Churners show wider spread and heavier upper-tail skew, suggesting intense binge consumption before canceling.

Hours Watched

Median activity for Rapid Churners is similar to Retained Subscribers — high usage in month 1 does not guarantee long-term loyalty.

Takeaway

Churn risk can be tied to usage patterns and variability, not just total activity volume.

Engagement metrics normalized by active months. Synthetic values for illustration.

Acquisition and Billing

Where subscribers come from shapes how long they stay.

Refund Patterns Reveal Onboarding Friction

ACQUISITION

Refunds are small in volume but concentrated in specific channels — pointing to acquisition friction.

~2.8%
of subscribers refunded

~830
refund requests in cohort

1 channel
drives majority of refunds

Internal billing channel A

Accounts for the majority of refund requests across the cohort.

TV signup platform B

Shows noticeably higher refund rates than other signup platforms.

Channel-specific friction

Concentration suggests onboarding issues are platform-specific, not systemic.

Takeaway: Small in volume, but refunds point to channel-specific onboarding issues — fixing these could prevent early churn and lost trust.

Channel labels anonymized. Synthetic values for illustration.

Churn Risk Varies by Acquisition Channel

ACQUISITION

Some billers and platforms drive higher churn than others.

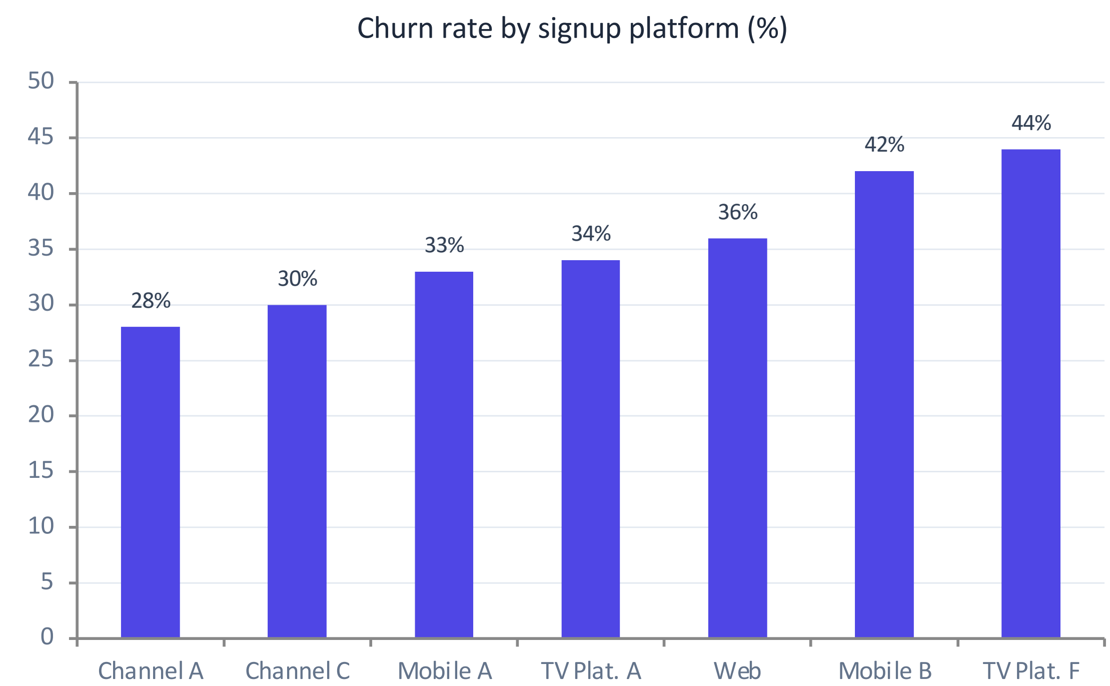

Acquisition Billers

~9%

higher churn

External-B and Internal-B churn ~9 percentage points higher than other billers.

Signup Platforms

>40%

churn rate

Mobile-B and TV Platform-F have highest churn vs. ~34% average.

Channel labels anonymized. Synthetic values for illustration.

Survival Analysis

How retention changes over time, and where the cliffs are.

Early Retention Defines Long-Term Survival

SURVIVAL

Most churn happens in the first month. Retention stabilizes after month 2.

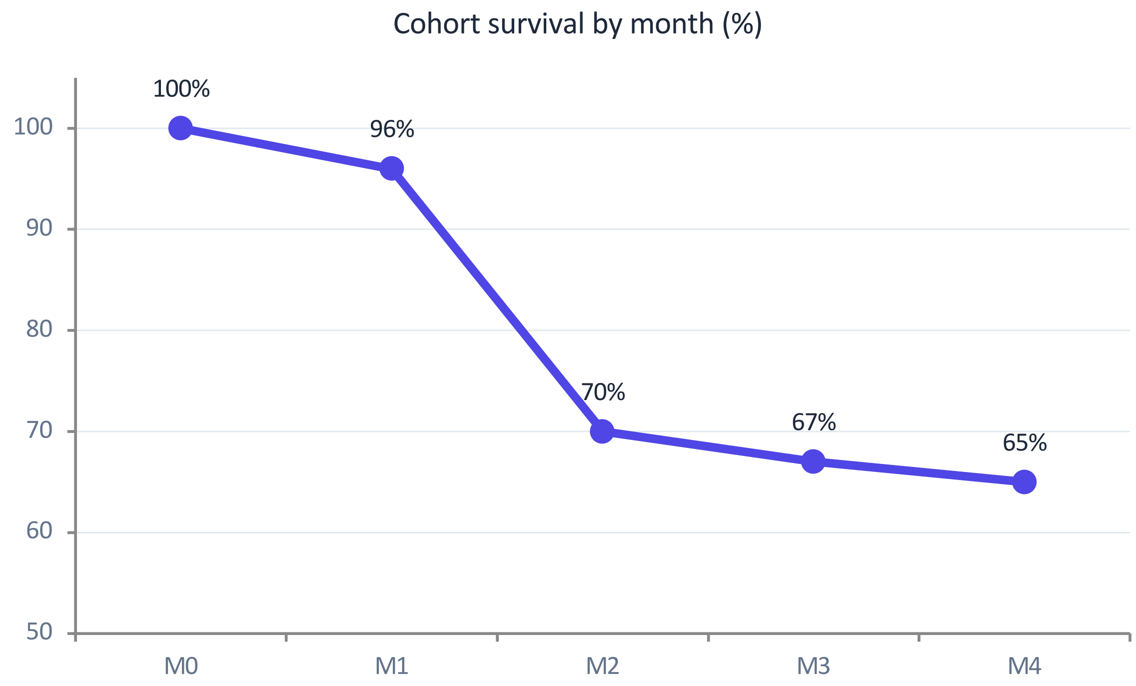

Steepest drop after month 1

Showing how critical the first month is to long-term retention.

Decline flattens after month 2

Subscribers who stay past this point are far more likely to remain.

Strategic implication

Strong onboarding and early engagement strategies are essential to capture long-term value.

Note: Final month data may be incomplete due to observation window cutoff. Synthetic values for illustration.

Survival by Segment: Where Churn Concentrates

SURVIVAL

Segment-level survival curves confirm that acquisition channel and engagement pattern drive retention.

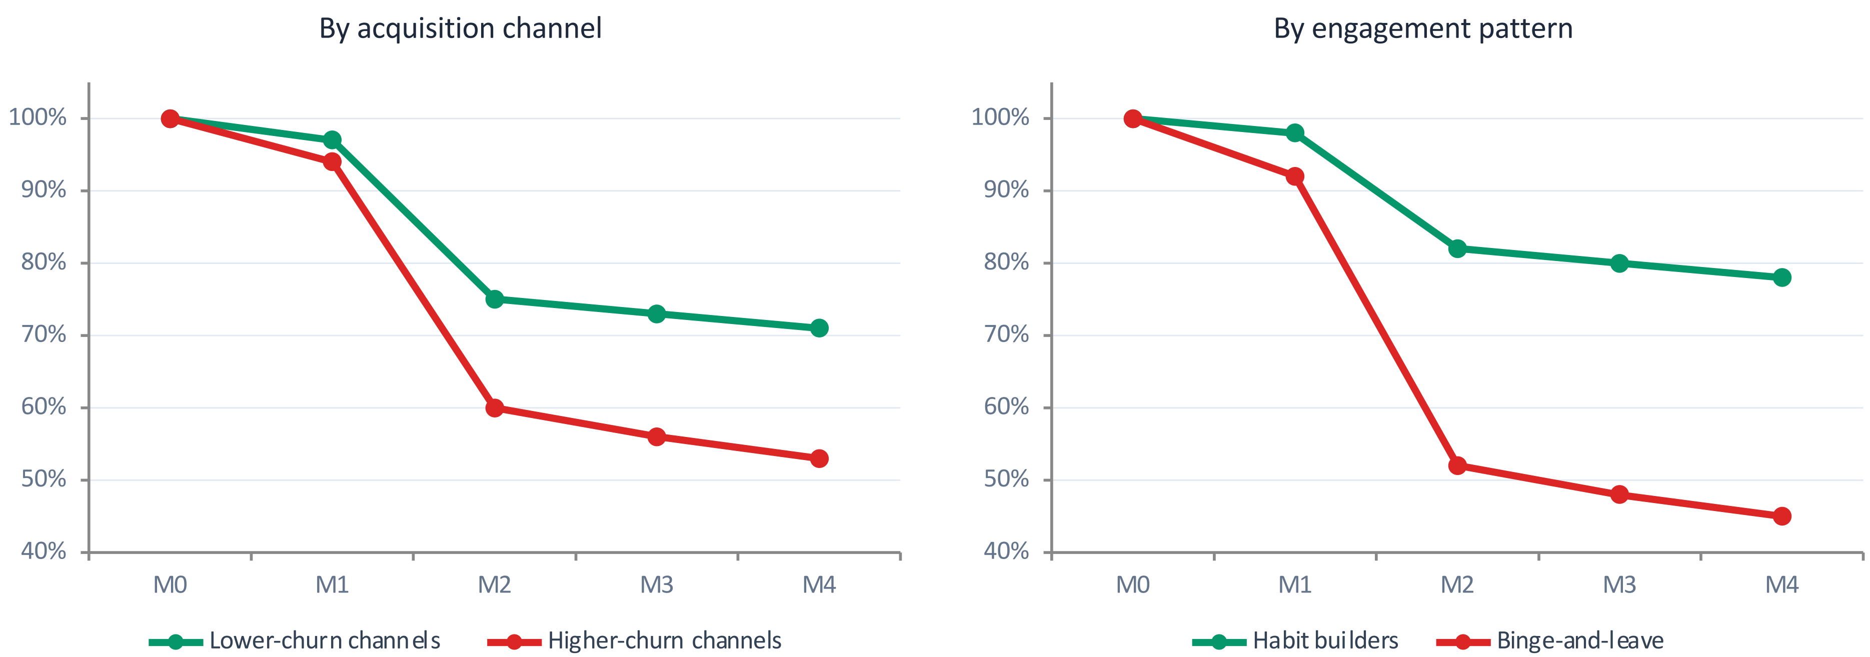

Across both cuts, survival diverges quickly in the first 30 days and tracks together afterward.

Note: Final month data may be incomplete. Synthetic values for illustration.

Survival by Tenure Tier and Billing Method

SURVIVAL

Further segmentation surfaces secondary risk factors for early churn.

Patterns are consistent: early-month divergence, then stabilization. Acquisition quality remains the dominant signal.

Note: Final month data may be incomplete. Synthetic values for illustration.

Summary: What Drives Churn

SYNTHESIS

Not all engagement equals loyalty. Early behavior pattern and acquisition channel drive churn.

01

"Binge and Leave" Pattern

Rapid churners show higher first-month stream starts and slightly higher hours — they binge what they came for, then cancel.

02

Habit Builders Retain

Longer-tenure subscribers show steadier, sustained usage patterns that align with higher retention rates.

03

High-Churn Channels

Specific billers and signup platforms exhibit above-average churn, suggesting acquisition path predicts tenure.

04

Refunds Indicate Friction

Refunds are small in volume but concentrated in specific billers and platforms — pointing to onboarding issues.

05

Survival Insight

Most attrition occurs after month 1, then stabilizes. The first month is the critical retention window.

Takeaway: Churn risk is concentrated in the first month. Subscribers who binge in month 1, or signed up via higher-churn channels, are least likely to stick.

03

Predictive Model

Results, limitations, and operational implications

Predictive Model Results

PREDICTIVE

The model gives some ability to flag early churners, but overall accuracy is limited due to data constraints.

~55-56%

Overall Accuracy

Slightly better than a coin flip

~67%

Precision (when flagged)

When model predicts churn, it is right ~2/3 of the time

~54%

Recall on Churners

Of those who churned, the model catches just over half

INTERPRETATION

The model finds some useful patterns but its predictive power is limited. Two-thirds precision means flagged subscribers are real risks worth contacting, but only catching half of actual churners means many slip through. The current data window and feature set don't separate "binge and leave" users from "habit builders" reliably.

Implications and Next Steps

STRATEGY

More data is needed to make churn predictions actionable, but actionable insights exist today.

WHAT WORKS TODAY

Channel-level risk is actionable now

Higher churn risks from certain acquisition billers and signup platforms can guide onboarding and retention focus immediately, without waiting for model improvement.

CURRENT LIMITATIONS

Data gaps cap model lift

Without weekly engagement data to capture early viewing behavior, the model can't reliably separate "binge and leave" from "habit builders." More demographic features would help.

FUTURE IMPROVEMENTS

Expand data, validate cohort

Add early engagement signals, expand timeframe with historical data, and include other cohorts to determine whether observed patterns are recurring or cohort-specific.

BUSINESS IMPACT

Modest accuracy, real value

Even with current accuracy, the analysis identifies higher-risk acquisition channels and platforms that can guide where retention investment lands first.