GROWTH THAT COMPOUNDS

Retention should be ranked by profit, not churn alone

High churn risk is not customer value: rank retention by incremental effect, intervention cost, postcampaign cash flow, and campaign size.

Sinan Isoglu, MBA (Quantic)

Commercial growth leader, lecturer and doctoral researcher

PUBLISHED
1 September 2026

READING TIME
7 min

WORDS
1,625

https://isoglu.com/insights/journal/retention-should-be-ranked-by-profit-not-churn-alone/

CONTENTS

MANAGEMENT SUMMARY

A customer can be likely to churn and still be a poor retention target if the intervention is costly or the relationship produces little incremental cash flow. Lemmens and Gupta define a profit-based loss function that ranks customers by the financial impact of an intervention, including its effect on churn, postcampaign cash flows, and intervention cost. Two field experiments report more profitable campaigns than competing targeting models and show how campaign size can be optimized. This article separates churn risk, response, incremental effect, cost, and profit lift in a synthetic review card. It does not rank a current customer base, provide a universal churn threshold, or evaluate private campaign data.

Keywords: Profit-based retention · Retention economics · Customer profitability · Churn targeting

A customer can be likely to churn and still be a poor retention target.

The short answer is that retention targeting is an intervention-allocation decision, not only a risk prediction. A churn score estimates a risk. It does not by itself show whether an offer changes that risk, what the intervention costs, or how much postcampaign cash flow remains after the intervention.

Lemmens and Gupta define a profit-based loss function for this problem. It ranks customers by the incremental financial impact of intervening, including changes in churn, postcampaign cash flows, and the cost of the intervention. Two field experiments report more profitable campaigns than competing targeting models and provide a way to optimize campaign size. The evidence supports a profit-first review boundary, not a universal targeting formula.

Two adjacent studies sharpen the boundary. Gupta et al. (2004) show that customer-based valuation depends on retention, contribution margin, customer count, acquisition, and discounting assumptions. Lariviere and Van den Poel (2005) model next purchase, partial defection, and profitability evolution as separate outcomes. Those findings do not replace profit-based targeting. They explain why a retention score must state its outcome and value horizon before it is used for intervention allocation.

Why is statistical churn probability an insufficient basis for customer retention spend?

Imagine a retention list sorted from the highest predicted churn probability to the lowest. The first customer may be valuable, or may be expensive to serve, unlikely to respond, or already close to leaving for reasons the offer cannot change. A high score identifies a problem to investigate. It does not finish the allocation decision.

Keep these objects separate:

Table 1 Why is statistical churn probability an insufficient basis for customer retention spend?

Object	Question	What it does not answer
Churn risk	How likely is the customer to leave under the stated prediction frame?	Whether an offer changes the risk
Response	How likely is the customer to react to an offer?	Whether the reaction creates profit
Incremental effect	What changes because the intervention is made rather than withheld?	Whether the change covers its cost
Intervention cost	What does it cost to make and deliver the offer?	Whether the customer would have stayed without it
Postcampaign cash flow	What financial flow remains after the intervention under the declared horizon?	Whether the customer is valuable in every future period
Profit lift	What incremental value remains after cost?	Whether the estimate transfers to another setting

Table from this essay. Sources and interpretation are given in the article.

The distinctions are not academic decoration. If a campaign ranks risk while ignoring intervention effect and cost, it can spend scarce capacity on customers who were going to stay, could not be influenced, or would not repay the offer.

Table 1 The profit-first retention ranking card

Synthetic customer	Churn risk	Expected incremental cash flow	Intervention cost	Profit lift after cost	Decision
A	High	42	4	38	Target
B	High	9	5	4	Hold for review
C	Medium	30	3	27	Target
D	Low	16	12	4	Hold for review

Author's synthetic decision framework grounded in Lemmens and Gupta (2020), Lariviere and Van den Poel (2005), and Gupta et al. (2004). Values are illustrative and do not describe observed customers or a live campaign.

The table is not a model. It is a reminder that a risk ranking and a campaign ranking answer different questions. Customer A and customer B have the same risk category but different expected economics. Customer C has lower risk than A but a larger synthetic profit lift. Customer D may be worth a different treatment, but the listed intervention does not justify a target decision under this made-up boundary.

How does a profit-based loss function transform retention campaign targeting?

Lemmens and Gupta's contribution begins with the objective. Conventional targeting can rank customers by predicted churn probability or by responsiveness to a retention offer. Both can be useful inputs. Neither aligns the prediction objective with campaign profit when the cost and financial consequences of an intervention vary across customers.

The profit-based loss function predicts the financial impact of the retention intervention for each customer. The ranking then uses the incremental effect on churn and postcampaign cash flows after the intervention cost. In other words, the object is not "who is most at risk?" It is "where does making this intervention, compared with not making it, create the most expected profit under the declared frame?"

That comparison requires a counterfactual boundary. The relevant quantity is incremental, not simply observed revenue after an offer. A customer may renew after receiving an offer because the offer helped, because the customer would have renewed anyway, or because some other event changed the relationship. The source's approach is designed around the financial impact of the intervention, while this article keeps the decision rule visible without reconstructing a current campaign.

Why is marginal profit lift distinct from cumulative customer lifetime value?

Customer Lifetime Value describes a longer-run customer-economics construct under a declared revenue, cost, retention, discounting, and time boundary. Gupta et al. (2004) make those valuation assumptions explicit, including the role of customer count and contribution margin. Profit-based retention targeting asks a narrower question about an intervention and its incremental financial effect. The two can inform one another, but they are not interchangeable.

A high-lifetime-value customer may deserve attention for many reasons and still be a poor target for one specific offer. A lower-value customer may be an efficient target if the intervention is cheap and changes the future cash flow enough to cover its cost. The source does not turn that observation into a universal customer hierarchy. It changes which variables belong in the campaign-ranking problem.

The same distinction applies to response. Lariviere and Van den Poel (2005) show why next purchase, partial defection, and profitability evolution should not be treated as one outcome. A customer can respond to an offer by opening an email, accepting a discount, or changing a service choice. A response is not automatically a profitable incremental effect. The campaign has to define the outcome and the cost of producing it.

How should commercial finance determine optimal retention campaign size?

Even a profit-ranked list does not decide how many customers to contact. The next customer in the ranking may add less expected profit than the operational cost of extending the campaign. The available capacity, offer inventory, contact rules, and expected interference between customers can also change the boundary.

Lemmens and Gupta report that their approach can optimize campaign size and that two field experiments find more profitable campaigns than competing targeting models. Those are source-reported results in the settings studied. They are not a universal rule that every campaign should stop at the same threshold or that a profit score can be transferred without recalibration.

Use a campaign-size gate:

1. declare the intervention and its comparison condition;
2. estimate incremental effect rather than reading response as impact;
3. include the intervention cost and the postcampaign cash-flow horizon;
4. rank the eligible customers under the same value boundary;
5. identify the operational cost of contacting the next group;
6. stop or extend the campaign based on the declared decision rule;
7. record which assumptions would change that rule.

The last step matters because a campaign can be profitable in expectation and still be wrong for a different horizon, service constraint, or customer population.

How can marketing teams construct an auditable profit-ranked retention scorecard?

Before using a profit-first ranking, write down five fields:

Table 3 How can marketing teams construct an auditable profit-ranked retention scorecard?

Review field	Required question
Intervention	What exactly is being offered, changed, or withheld?
Incrementality	Compared with what would have happened without the intervention?
Economics	Which revenue, cost, and cash-flow items enter the calculation?
Capacity	How many contacts, offers, or service actions can be delivered?
Validation	What outcome and horizon will tell the team whether the ranking worked?

Table from this essay. Sources and interpretation are given in the article.

These fields do not produce a score by themselves. They make a score challengeable. A reader can ask whether the intervention changes churn, whether cost is included, whether the horizon is long enough, and whether the result is observed in the same setting as the model.

Which three conventional customer retention assumptions should leadership reject?

First, do not say “target customers most likely to churn” when the decision is actually to maximize profit. Risk can be one feature of the ranking, but it is not the objective by definition.

Second, do not say “the campaign created the retained revenue” when the evidence is only a postcampaign observation. Incrementality needs a comparison boundary.

Third, do not say “the model finds the best customers” without naming the intervention, cost, cash-flow horizon, and outcome. Best for what decision is the only useful version of that sentence.

For adjacent decisions, compare the renewal evidence review with the retention-rate measurement boundary.

Where are the empirical boundaries of profit-based retention modeling?

Lemmens and Gupta provide a profit-based loss function, an incremental intervention-ranking objective, two field experiments, and a campaign-size implication. Gupta et al. provide a customer-value assumption boundary, while Lariviere and Van den Poel provide an outcome-specific prediction boundary. None of the three studies evaluates the owner’s customer base or establishes a universal churn threshold. The synthetic card is an author-owned translation. It shows what a profit-first review must keep visible before a retention list becomes a spending decision.

Lemmens, A., & Gupta, S. (2020). Managing churn to maximize profits. Marketing Science, 39(5), 956-973. DOI

Gupta, S., Lehmann, D. R., & Stuart, J. A. (2004). Valuing customers. Journal of Marketing Research, 41(1), 7-18. DOI

Lariviere, B., & Van den Poel, D. (2005). Predicting customer retention and profitability by using random forests and regression forests techniques. Expert Systems with Applications, 29(2), 398-405. DOI

END OF ESSAY

HOW TO CITE THIS ESSAY

Isoglu, S. (2026, September 1). Retention should be ranked by profit, not churn alone. Sinan Isoglu. https://isoglu.com/insights/journal/retention-should-be-ranked-by-profit-not-churn-alone/

KEEP READING

Growth that compounds

Long customer relationships can be low-profit

https://isoglu.com/insights/journal/long-customer-relationships-can-be-low-profit/

Growth that compounds

What is customer churn?

https://isoglu.com/insights/journal/what-is-customer-churn/

Growth that compounds

What is customer lifetime value?

https://isoglu.com/insights/journal/what-is-clv/

Sinan Isoglu, MBA (Quantic)
Commercial growth leader, lecturer and doctoral researcher

Online version: https://isoglu.com/insights/journal/retention-should-be-ranked-by-profit-not-churn-alone/