FINTECH & NEOBANKS · SUBSCRIPTION TIER UPGRADE PROMPT TIMING

Lyra

Subscription tier upgrade prompt timing

CUSTOMER
Series-B neobank · ~$18M ARR · MENA + EU

PERIOD
Jan – Mar 2026 (90 days)

SAMPLE
482,711 eligible free-tier accounts

EXPERIMENT AUDITED
Premium upgrade prompt timing · week-4 (control) vs week-2 (variant)

FINTECH & NEOBANKS

Lyra · Subscription tier upgrade prompt timing

1 · WHAT THEIR TEAM REPORTED

METRIC	WEEK-4 (CONTROL)	WEEK-2 (VARIANT)
Upgrade rate	4.1%	4.4%
p-value	—	0.02

> *Team report: "winner — ship week-2." Variant shipped to 100 % of free-tier users.*

2 · OUR RE-ANALYSIS · DOUBLY-ROBUST + PER-SEGMENT CATE

COHORT	DR ESTIMATE	95% CI	ESS	VERDICT	$ IMPACT / YR
All free-tier accounts	+6.4% rel.	[+3.0, +9.8]	0.58	positive — larger than t-test	—
Balance high · spend regular	**+11.7% rel.**	**[+7.2, +16.3]**	**0.51**	clear positive uplift	**+$0.6M (ship)**
Balance high · spend irregular	+5.1% rel.	[+1.0, +9.3]	0.48	small positive	+$0.2M (ship)
Balance low · spend regular	+2.4% rel.	[−2.1, +6.9]	0.42	inconclusive	—
Balance low · spend irregular	**−8.2% rel.**	**[−13.5, −2.9]**	**0.36**	clear negative — prompted into churn	**+$0.6M (suppress)**
New accounts (<30d tenure)	−3.1% rel.	[−7.2, +1.0]	0.31	overlap-limited; flag	— (re-test)

3 · THE HIDDEN COHORT EFFECT

The negative-uplift cell ({balance-low × spend-irregular}, ~14 % of eligible base) is being prompted into downgrade or cancellation. Suppressing the prompt for that cohort would lift **net** upgrade rate from +6.4 % to an estimated **+11.7 %**.

4 · WHAT WE'D RECOMMEND

Personalise the prompt. Send to positive-uplift cohorts (~58 % of base), suppress for the negative cell.

PROJECTED ANNUALISED IMPACT

+$1.4M / yr ARR · retention preserved on the irregular-spend cohort

Estimates use 1,000-bootstrap doubly-robust evaluation against the customer's logged data with propensities reconstructed from the experiment configuration. ESS < n/10 cells are flagged overlap-limited and not used in the recommendation. The customer can reproduce these numbers on their own logs using `offpolicy.py` (MIT-licensed). Numbers in this report are fictive, generated as a worked example — not derived from real customer data.